An Interop Publication

CONNEXIONS®

The Interoperability Report

November 1996

Volume 10, No. 11

ConneXions — The Interoperability Report tracks current and emerging standards and technologies within the computer and communications industry.

In this issue:

ConneXions is published monthly by Interop Company, a division of SOFTBANK Exposition and Conference Company, 303 Vintage Park Drive, Foster City, California, 94404–1138, USA.
Phone: +1 (415) 578-6900
Fax: +1 (415) 525-0194
E-mail: connexions@interop.com

Subscription hotline: 1-800-575-5717
or +1 610-892-1959

Copyright © 1996 by Interop Company. Quotation with attribution encouraged.

ConneXions—The Interoperability Report and the *ConneXions* logo are registered trademarks of Interop Company.

ISSN 0894–5926

From the Editor

After nearly 10 years, 117 issues and more than 3,600 printed pages, *ConneXions—The Interoperability Report* will retire as a print publication at the end of 1996. A new Web-based publication, *ConneXions Online,* will emerge in early 1997. We will continue to publish in-depth technical tutorials on all aspects of networking in this new medium. In addition, we plan to make our entire collection of back issues available on the Web, thus creating a large digital library of networking reference material. Access to *ConneXions Online* and its collection of reference material will be free of charge.

Needless to say, the creation of this digital library and the implementation of a suitable search front-end will not happen overnight, but will evolve over a period of time. Key players in this project will be our online development group WebSource assisted by myself as the editor, and we hope a number of volunteers interested in such a project. The collection will have links to other sources of information such as the Request For Comments (RFC) document series, online glossaries and so on. Special advisor on this project is Jack Kessler, kessler@well.sf.ca.us, Internet Trainer and Consultant, who is an expert on digital libraries.

I would like to take this opportunity to thank all the people who have contributed to *ConneXions* since it started publication in March 1987. Several hundred authors have shared their expertise with our readers. Their work is what makes *ConneXions* possible. Special thanks goes to the Editorial Advisory Board as well as the crew at Globe Printing Company. Countless readers have provided valuable feedback and suggestions. Last, but not least, the folks at Seybold Publications, our subscription agency, deserve a round of applause for all their customer service functions.

I hope you will continue to support *ConneXions Online* and help us make it a 21st century publication. Comments, suggestions, and above all *articles* are welcome as always. HTML is not the only acceptable format, but it would obviously help ;–). Send your letters and articles to connexions@interop.com.

This month, we take a look ahead to the year 2000 and discover what may happen to computer systems unless proper precautions are taken. We also look at security in large networks and learn about ways to improve ATM cell processing in SDH multiplexers. Finally, a brief tutorial on the two emerging standards for 100Mbps Ethernet, 100VG-AnyLAN and 100-BaseT. We plan to publish some interesting test results on these two technologies in our final issue.

The Millennium Bomb

by Ian Hugo

Introduction

You have probably heard of the Year 2000 problem. Much software uses just two digits to represent the year—and this software will fail when the clock ticks over into the next century, and the date appears to go to zero.

It is popularly believed that only mainframes and legacy COBOL, PL/1 and Assembler code are affected by the year 2000 problem. In fact, this is not the case.

We know of only one case where the IS department went as far as auditing the distributed environment as well as the mainframe data centre. And in that case, the distributed system turned out to be the larger problem. In the next few pages we'll disabuse you of the myths, explain how you may be affected and suggest what you might do about it. We will focus on systems and service management.

Firstly, though, we'll tell you why your organisation should be worried in general.

Why worry?

The principal reason to be worried is that most organisations apparently aren't. They ought to be in a state of at least mild panic. There's a rapidly growing body of evidence suggesting that, if you're not panicking, you don't understand the problem.

Mike Kerford-Byrnes, an independent consultant specialising in the problem, states: "I have been talking about this since 1984 and still many people don't want to listen. One mistake I think they make is to assume that it's just a matter of changing a few dates; in fact, that's the smallest part of the problem. By far the largest parts are establishing exactly what software is affected and how, working out a strategy to resolve the problem and then testing all the corrections after they have been made. That will account for around 80 percent of the work and, even if you think you're clear, you are still going to have to do the testing."

A big part of the work is an impact analysis. If that sounds trivial, it isn't. Leaving aside the question of whether you have all source code for aged applications (you won't have), there's the need for a complete, accurate and up-to-date inventory of all software you use. You won't have that either, and that's before you even start doing an impact analysis.

Most sites active on the problem have done a cascade of impact analyses, using information gleaned at higher levels to justify more detailed analyses. All the available evidence is that early estimates are always revised, upwards. Industry watchers Xephon have been conducting annual surveys and, this year, one finding is that companies working on the problem last year have revised their estimates upwards this year by a factor of between 3 and 6. They also found that the number of sites who thought they were clear of the problem has decreased over the last year.

What size of problem are we talking about? For a mid-range large organisation, with around 3,500–4,000 programs to review, the problem translates into about 100 man-years of effort and a spend of around £6 million. The spend increases in inverse proportion to the time available; The Gartner Group has revised its cost estimates upwards by a factor of 5 over the past two years. What benefit do you get for this spend? Essentially nothing; programs simply run as they did before.

That, in fact, is really where the initial problem lies. How does anyone persuade their Board to spend several million pounds for no business benefit? That's the real worry; here's some more detail.

Myths in general

Firstly, the Year 2000 problem has almost nothing whatsoever to do with specific programming languages. It is all about date representations and the logic applied to them. So, problems can occur just as easily with **C**, **C++**, Basic and Visual Basic programs (as well as Easytrieve, Mark IV, Dibol, BusinessBasic, Nomad, etc) as with programs written in any other language.

The PC problem

There's a belief in some circles that Year 2000 problems don't affect PCs or small systems that primarily use packages. This is *not* true. On your own (personal) PC, you can conduct a simple test right now. First make sure that you have no software loaded that you can't afford to lose and which you can't re-install. Set the system date to 31 December 1999 and the time to 23:58 hrs; then switch the machine off. Wait at least three minutes and then switch your machine on. What do you expect to find?

There's a good chance that the machine will reboot with a new date of 4 January 1980; or it may revert to the 1 January 1900. Those are common results, although a range of other bizarre results also occur. There's also a chance that your system will simply refuse to reboot altogether. Any software which you had the right to use for less than four years will obviously refuse to work and, if that includes your system software, you have a dead machine.

Exactly what will happen depends on the precise hardware configuration you are using and the operating system and BIOS installed. Windows 95 appears to be relatively clear of problems on most machines; DOS and Windows 3.1 will generally produce problems of one sort or another.

However, as US industry commentator Peter de Jager, who suggests this test, points out: "This is just a test; on the 1st of January 2000 it will be painfully real. You may be aware now and able to correct the fault but, on the 31st of December 1999, some 80 million PCs will be switched off; who ensures that all 80 million users are aware and able to make the necessary corrections in time?"

In fact, if you want to know in advance how your PC will react to this test, without risking it, you can contact Geoff May via e-mail at `geoffmay@enternet.com.au`, who is keeping a log of how different hardware/operating system/BIOS combinations react to the problem.

And that was just the system—what about applications? Ken Lindup, spokesperson on Year 2000 issues for the *Institute of Internal Auditors,* says: "The problem is massive. Even at the PC level, there are lots and lots of bespoke Dbase applications in use that will cause problems." Even then, we're still just considering individual PCs. Add in PC networks, servers, mainframes and dependencies on external systems, through supply chains, EDI, credit clearing, etc, and the dimensions of the problem become potentially awesome.

A little detail ...

It is worth setting out the problem in some detail here. Just in case you're not already familiar with it, the problem derives from two points: firstly, assumptions made in the logic of virtually all routines which access date, that the numerical value of the date always increases as time goes by; secondly, the fact that dates are nearly always held in the form YYMMDD, with just 2 characters for the year. This poses no problem until 991231 (31 December 1999), which suddenly becomes 000101 (1 January 2000).

continued on next page

The Millennium Bomb *(continued)*

It is elementary to see that the later date has a lower value than the earlier date, thus destroying the logic of programs of all kinds.

The effects of the false logic range from the bizarre, as in the PC test, through the hilarious, to the catastrophic. We can already see the kind of thing that can happen when people's ages are encoded in two digits—two American centegenarians (aged 104) have recently been invited to start school. More seriously, newly manufactured goods have been barcoded as past their expiry date and prevented only by chance human intervention from being scrapped. Potentially, the same can happen to your backup files and archived data.

The problem is progressive for applications that look ahead a number of years, and already tragicomic occurrences are being reported. Despite this, numerous organisations seem still to be taking a relaxed view of the problem, regarding it as a minor hiccup that they will deal with later.

Costs

Estimates of the size of the problem depend necessarily on guesswork to a large degree, since no large, detailed surveys have been done and anyway the details can't be known until most organisations start impact analyses. The Gartner Group has been quoted as estimating the potential cost worldwide as $600 billion. Ian Baker, IBM's Year 2000 Services Manager, thinks it may be nearer $400 billion. He says: "I think about a third of this cost will be incurred in Europe, say $125 billion..."

The misconception that the problem applies only to mainframes and legacy COBOL, PL/1 and Assembler stems from the fact that these are the worst cases; but the problem is far from being confined to them alone, as we have already illustrated with PCs. In fact, the programming language used is irrelevant; any code written in **C** and **C++** that doesn't take account of the new millennium will fail just as surely as 20-year old COBOL programs.

And, incidentally, the Year 2000 is also a leap year, a routine occurrence every 4-years since computers were invented. Yet systems at Papworth hospital and the UK Meteorological Office were reported to have failed on February 29th this year because they couldn't cope with that.

Finally, there is the idea that, if you use packages, it's a problem for the package supplier rather than you. That may be true, provided you are running a recent release of the packages, haven't modified the code in any way and don't pass files between packages from different suppliers. If you're running old releases of a package, you'll probably have to pay for an upgrade (plus, of course, upgrades to any dependent software not supported by the new release). If you've modified packages, Ian Baker suggests you may find it impossible to retrofit the modifications to new releases. And if you pass files between packages, then you need to know that the corrections adopted by the package suppliers are mutually coherent. For instance, one supplier may resolve the problem by adding two digits to the year representation while another decides to keep to two digits for the year and code round the problem.

Everyone is affected

So, everyone is affected, or will be, in one way or another. If it is difficult to understand why there is so much more work to do than just making the date changes, consider these points.

Just to start, you have to review all the software you are using.

Do you have a complete, up-to-date and accurate software inventory, even for your mainframes and servers let alone your PCs? Probably no organisation has. Secondly, do you still have source code and documentation for all programs developed or modified in-house across all platforms? Again, probably no organisation has that either.

When you've made a pass at finding all your software, you then start looking for dates, not just in applications but also in any system management or other utilities developed or modified in-house. (Incidentally, you also need to get from all your software suppliers what solutions they propose to adopt and when compliant releases will be available.) Finding dates may not be as easy as you expect.

ICL consultant Janet Pavelin points out: "Dates are scattered through tables, records and databases, and in a variety of formats and representations (sequence of day, month, year); they can be numeric (signed binary, packed, display) or alphanumeric. Reports and screens often have "19" hard-coded as the year, so casual inspection is not a way to check. Also, two-digit dates can appear as database keys sometimes embedded in serial numbers and other identifiers that may be sorted, which makes them hard to detect."

Anyway, once you have sorted out that little problem, you then have to decide on a strategy for correcting the potential errors. Very broadly, there are two possible strategies: one is to add two characters to representation of the year, the other to code round the problem in some way. The former solution solves the problem forever (potentially; date counters in some hardware/OS systems can still overflow), the latter solution is generally favoured as being less disruptive. However, the code used to resolve the problem then remains as a potential trap for future maintenance staff.

If all this seems a somewhat elaborate explanation of what is involved, you won't want to know that this accounts for only around 60 percent of the effort. You still have to do the testing and testing is estimated by Mike Kerford-Byrnes and others to account for some 40 percent of the work.

If, again, this figure sounds outrageous, consider the following three points. Firstly, if you want to be absolutely sure you're clear (and you might care to think about your BS5750/ES9000 quality certification in this context), you have to test all your software, including any that you believe to be unaffected. Secondly, testing of individual programs and affected modules is a small part of the task. A full regression test, of applications, files, databases and all relevant system software (all in supposedly compliant versions) is required. You'll probably need a separate system for this anyway but it's worth noting that testing requires the system clock to be advanced to 01.01.2000, which you certainly wouldn't do on a production machine. You could test on a development/test machine by inserting code which intercepts calls for the system date and modifying the result.

Even so, you'll probably want to do the testing over a weekend and, Mike Kerford-Byrnes points out: "There are fewer than 140 weekends left between now and the 1st of January 1999. You'll need 1999 to clear up any final problems and to test connections with external systems." Apart from which, if you haven't yet started on an impact analysis, you're a long way from even beginning testing.

In fact, every organisation will become increasingly aware of the problem because, as we pointed out earlier, the problem is progressive. Most applications look ahead to some extent, the large majority for at least a year.

continued on next page

The Millennium Bomb *(continued)*

The Gartner Group estimates that 90 percent of applications will be affected by 1999. The danger is that, if you don't put a formal strategy in place very quickly, different ad hoc solutions to affected applications will compound the problem.

Remember, next time you hear Big Ben toll in the New Year, the words of poet John Donne: "Send not to know for whom the bell tolls; it tolls for thee."

IAN HUGO is editor of *Millennium Watch,* a newsletter produced in association with the UK government-initiated Taskforce 2000. He runs VISUAL, a cross-vendor user group focussed on new methods tools and techniques. He undertakes technical marketing assignments for major IT corporations. He can be contacted on `ihugo@hassocs.netkonnect.co.uk`

Year 2000: The Human Problem

by Peter Judge

Introduction

The article above makes it clear that the Year 2000 is not just a problem for mainframes, but one for distributed systems as well. This article is to illustrate that the problem may go beyond that.

As medieval Europe approached the year 1000, millennial racketeers persuaded the gullible to leave their homes and farms, to prepare for the new era of prosperity which was about to dawn. Now, as the world turns towards the year 2000, is the Year 2000 problem another racket? Or is it real? I believe it is real, but that users should take care which version of it they believe.

With three and half years to go, the Year 2000 is finally being taken seriously by the industry, and even by politicians. In the UK, for instance, the opposition Labour Party is already warning that serious investment will be required to solve the problem in government systems. A general election next year might land the problem—and the problem of finding an undetermined amount of money to solve it—entirely on the shoulders of a new Labour government.

Meanwhile, the UK government is planning a publicity campaign to get businesses to act on the issue, but it appears that no government departments have estimated their own liability yet.

The US public sector could lose $30 billion on this, and the House of Representatives has heard testimony about it.

The more dramatic visions of the Year 2000 are apocalyptic, with looting in the streets and a (hopefully temporary) breakdown of some of the infrastructure of society. The source of these visions is easy to see: airlines and other transport, banks, power stations, communications, welfare systems and emergency services all depend on computers. If those computers fail, will the service go down too?

This, combined with the high expectations and lowered inhibitions created by the millennium could lead to civil disorder among the disadvantaged.

A financial doom scenario might emerge from the failure of credit card systems. If those systems fail, or if there is a widespread expectation that they will, the middle classes of America and the developed world might all draw out enough cash to cover them for the first weeks of the year. The result could be that banks crash, and a repeat of the Stock Market crash of the 1920s.

Is it a racket?

Cynics might argue that this kind of vision is self-serving—Year 2000 services are being sold by playing on people's fear of the unknown. This argument says the Year 2000 bandwagon is a racket.

"The Year 2000 problem has three advantages that help to make it a perfect racket," warns Nicholas Zvegintzov, in an article in *American Programmer*. "First, it has a basis in reality. Years stored as two decimal digits *will* overflow in the year 2000. Second, it is a software problem—perhaps the only software problem—that lay people can easily understand. Many have watched an automobile odometer overflow from 999 to 000; many have seen forms printed with two few spaces for a date; many can see the potential problems. Writers and readers who are deeply ignorant of the nature and functioning of software are equally deeply grateful to find an issue that they can understand. It is unique in its comprehensibility."

"Third," he says, "its association with the turn of the millennium plays into superstition." He claims no technology since alchemy has excited the same level of superstition as computers. Incantations like "object-orientation" are seen as healing spells, and there is a superstition that children know more about IT than adults.

Solving the Year 2000 problem is an exercise for the software novice, says Mr. Svegintzov, who has edited *Software Maintenance News,* and serves on the Program Committee of the International Conference on Software Maintenance. "Most real world software problems are much harder; they are problems in which neither the context, the symptoms, the evidence, or the treatment are so plain. If an organization cannot handle the Year 2000 problem, it is a dangerous organization indeed; avoid it!"

According to his view, software professionals are colluding with the panic as a means to get funds for software problems which might otherwise be ignored or misunderstood by management.

A real problem

I do not agree. Certainly, individual Year 2000 problems are simple to solve, but the problem is dealing with large organisations which have:

- Interacting systems,
- Undocumented or altered code, and
- A big volume of changes to be made and tested.

And some problems may be hidden—for example, suppose a supplier used a hidden routine to disable software whose license had expired—and used two digit dates? And what if your automated back-up system (which works on dated files) fails, in such a way as to lose the data?

The creativity of programmers just adds to this problem. In some systems, extra information is coded into the date field by using the figures 99 and 00 for specific meanings, like "last entry" or "null record," or using them as interrupt markers. This software will fail in different and interesting ways.

The problem can get so convoluted it is difficult to know whether an organisation has solved it or not. Some large organisations are believed to have dealt with the problem as part of major systems upgrades in 1987. *Distributed Computing Directions* (from which this article is taken) has been assured that two major banks have fixed the problem—but has been shown evidence that at least one of them still uses two-digit dates.

continued on next page

Year 2000: The Human Problem *(continued)*

Benefits

There could be real benefits to solving the Year 2000 problem. Some issues can be fixed by software upgrades which have their own benefits in terms of new features and opportunities. For example, object-oriented code is less likely to have problems, and those problems will be easier to fix. And the other benefits of object orientation are well known.

The process of checking and testing will give the company the best software inventory it has ever had, and part of the project may invest in better management tools. Simply replacing the mainframe—and checking the new system carefully—could go a long way towards sorting the problem out. And, despite the current swing against open systems and downsizing, mainframe replacement still has a lot to recommend it.

But the irony is, that it is already too late for most large organisations to take these beneficial routes. As our companion article indicates, if businesses work quickly, they may just about get fixes applied to their existing systems. But for all but smaller applications it is too late to replace them.

Large-scale replacements cannot be done in a hurry, or with a fixed and immovable deadline. And to add to the irony, the expensive patches put on old software may extend its life, by siphoning off funds which could otherwise have gone towards replacing it.

It gives some measure of the scale of the problem that it has been suggested that all clocks be put back to 1980, or to 1950 to avoid the problem—though of course, software would still have to be changed to allow this time reversal to take place.

The problem may spread itself out. Some software will fail sooner—for example aircraft with preventive maintenance software are already refusing to fly because some of the dates when parts are due to fail have overflowed. Other systems such as generating plants may fail repeatedly up till the Millennium.

And on the day itself, the problem may strike later in the day, or several times in the day, when networked systems in different time zones move into the new year and overflow.

One result may be the failure of older organisations, and the survival of newer organisations, whose IS systems were developed more recently.

Leap year

And don't forget that 2000 is a leap year—and make sure your programmers all know it too. Although they divide by four, most new-century dates are not leap years, to correct for the real speed of the Earth's orbit round the Sun. This is quite well known, and I know of at least one mainframe package which has been delivered with the assumption that 2000 will not be a leap year.

The programmer in question was dismayed to hear that every four centuries a further correction is required—so the year 2000 will in fact be a leap year. The year 2000 is an exception to the exception!

A dry run

Talk to Unisys users—it is reported they have already had a go at this problem. According to Year 2000 expert Peter de Jager, Unisys 2200 machines showed faults on January 1 1996, when the 8th bit of the year field went to 1.

Getting the staff

Staffing a project as time-critical as this can be complicated. Most large organisations are giving their Year 2000 staff a bonus related to how successfully the organisation weathers the storm. Others have seen an opportunity in the sheer difficulty of finding staff to resolve problems in the holiday period at the beginning of the year 2000.

January 1, 2000 is a Saturday, and the following Monday is likely to be a holiday in most countries. Some parties may go on all week. Imagine the problem of finding a baby-sitter on Millennium Eve, and then consider the problem of getting good staff if you find you have a crisis on Millennium day. The problem of staffing will be added to by the difficulty of getting staff to your sites, and accommodating them, if they are required to work round the clock for the first days of the century.

As Mr. de Jager points out, can the emergency staff even do the work? Will the security system let them in, or will it see permits which have expired? Will the lifts work to get the staff into the data centre? And will the company PBX still work, or must they operate without phone contact? And what if the e-mail system loses all the messages?

Doing the work

The Year 2000 Web site (`http://www.year2000.com`) has a lot of useful resources on the subject. It include a checklist, by Serge Bouwens, which addresses many of the issues covered here, along with:

- Problems with your business partners.
- Problems imported during mergers and takeovers.
- Is the Year 2000 covered by service level agreements? Facilities management agreements may not lift the problem off your shoulders.
- Denial (you can't find the project without the budget to study it).

Join a user group

It makes sense, if you see the problem as real, to join a user group which is dealing with it. However, some companies are reluctant to do so. For reasons of competitive disadvantage they do not want to admit the problem in public. A user group can be a place to exchange information and experience, and get independent evaluation of the problems and the consultants' claims. It can provide management arguments for dealing with it effectively. The group can also focus attention, and ensure that the issue gets sufficient realistic and responsible publicity. And, if there is no local user group—you could start one.

In the end, the biggest beneficiaries of this will probably be the lawyers. Software created by third parties such as the big accountancy firms, will have been delivered to meet a contract which did not mention this problem. If problems arise, you can bet that both sides will argue over whether Year 2000 failure was included in the contract. Even contracts which do require an application to be millennium proof will be fought over—if the application fails due to a rogue date from another source.

The software maintenance community hopes that this issue will bring their discipline to the fore. Once this crisis is past, there are plenty of others to come. In the medium term, financial applications in Europe may have to face big changes in the event of a single currency.

And further in the future, our children will have to deal with a whole operating system failing, on 19 January 2038, when the UNIX date overflows.

[Ed.: Both Millennium articles are reprinted with permission from the June/July 1996 issue of *Distributed Computing Directions.* See `http://www.techapps.co.uk`]

PETER JUDGE edits and writes. He is editor of *Distributed Computing Directions*, *Internet Business* and the *Telecoms Newsline* email-zine. He is also writing, with Christopher Ogg, a report on the Internet and Intranets, due out this autumn from Cambridge Market Intelligence. He has degrees in Physics (Cambridge) and Fine Art (St Martin's, London). He has served on the programme committee for NetWorld+ Interop London. E-mail: `peter@pjudge.demon.co.uk`

Providing Security in a Large Network

by Craig A. Finseth,
InterTechnologies Group, State of Minnesota

Introduction

This article presents the approach to security taken by MNET. This approach enables MNET to offer some useful guarantees using existing technology.

Background

MNET is the data network for the State of Minnesota. It is operated by the InterTechnologies Group of the Department of Administration. Like the rest of its parent Group, MNET operates on a cost-recovery basis and receives no direct appropriations; we must compete on the open market for customers.

By law, MNET can serve any public sector body in the state. These potential customers include state agencies, county governments, city governments, K–12 schools, libraries, and the *Minnesota State Colleges and Universities* (MnSCU). As is apparent from this list, MNET's customer base has a wide variety of security needs.

MNET currently has twelve major hub sites around the state, which are linked together by multiple T1 circuits. Our Internet connection consists of multiple T1 circuits. We manage over 160 Cisco routers and have almost 300 customer connections. Together, our customers have on the order of 100,000 hosts connected to our network.

Security on those hosts ranges from very tight to very lax. They are managed by hundreds of separate organizations, each with their own agenda and priorities for security. Further, the networks that the hosts are attached to range from very secure to very insecure. In addition, each customer can have both dial-up connections and permanent "back doors" into other organizations' networks.

Over the past year, we have devised and implemented a security approach that allows us to provide tight security to those customers that want it and yet be an "open" provider to other customers.

The approach

MNET provides security in four different ways:

- Physical security
- Correct delivery of packets
- Certification of source addressing
- Customer-specified filters

Each will be discussed in turn.

Note that physical security applies to all MNET data traffic. The latter three only apply to TCP/IP traffic. TCP/IP is the only protocol sanctioned by MNET for use when security, privacy, or integrity are important.

Physical security

Our routers are located in either MNET-controlled space (i.e., our twelve major hub sites) or in space that is physically secure—at least as secure as the customer's own network.

Note that the customer networks that attach to our routers are *not* considered part of our network; each customer is responsible for the security of their own internal network.

The serial links from remote routers back to the rest of our network are all digital (56Kbps or faster) or fibre links. As such, they are not readily susceptible to wiretapping. However, even if a link were to be tapped, an intruder could only obtain data going to or from that particular site.

Correct delivery of packets

MNET routers are configured to not accept routing information from customers. Thus, regardless of what routing data a customer sends to us, it will be ignored by our routers. Hence, our routers will always deliver packets to the specified destinations.

Further, the bulk of the routers that we manage are "leaf" routers: those with only a single link back to the rest of the network. (These are also most apt to be the routers in the least physically secure locations.) For leaf routers, we use only static routing within our network. If an intruder were to compromise a leaf router, the rest of our network would ignore any routing information received from that router.

The only dynamic routing information exchanged is among the relatively few and well-protected core routers.

Certification of source addressing

First and foremost, all routers are configured to reject packets with the *IP Source Route* option turned on. This configuration helps our network resist one large class of possible attacks.

But further, all routers are configured to check the source addressing at every point at which customer traffic enters our network. With this filtering in place, we certify for every packet received from a customer that, if we were to deliver the packet to the *sending* address, we would deliver it back to the same interface that it was received by MNET from.

This certification is at the router interface level. That is, we can not and do not certify individual host addressing. Also, in the case where multiple networks are routed to a specific interface, we cannot certify that a host on one network is not masquerading as a host on another network off that same interface.

We apply similar filters to our Internet link, but in reverse. Essentially, we disallow any source addresses from the Internet that appear on any other customer connection.

Thus, for any particular IP address, it will be accepted on at most one MNET customer interface. (Addresses belonging to unassigned networks are not allowed in on any interface. Hence the "at most one.")

Customer specified filters

Each customer interface is also a place that we can (and do) provide protection for our customers: namely, *output access lists.* As part of our standard procedures, we require that our customers specify to us how to configure the outgoing filters on their interface (there is no default setting: we must receive a positive statement).

In general, we will accept any valid Cisco access list from a customer as their filter. However, most of our customers do not have advance Cisco training (:–), so we offer a number of "pre-packaged" filters.

The main choices are:

- None
- One-way
- One-way, except...
- Limited

continued on next page

Providing Security in a Large Network *(continued)*

"None" is simply no filtering, or wide open both ways. This choice is often made by customers who have their own firewall or who otherwise wish to take full responsibility for their own security.

"One-way" allows TCP connections to originate from the customer network to anywhere, but does not allow connections from anywhere else to be made to the customer network. (The only UDP traffic allowed is DNS, NTP, and SNMP requests and responses in the appropriate directions.) This type of filter is well-suited to those customers who have client-only computers.

"One-way, except..." is just like One-way, except that the customer specifies a small number of holes (e.g., for mail or WWW servers). This setting allows a customer to offer services while only having to worry about securing those services that they intend to offer.

"Limited" allows the customer to specify a list of addresses that can access this site. (In theory, it is a very large list, but in practice, the list is fairly short.) Because of the source address certification that we make, the filters can reliably distinguish among customers.

These choices can be combined in a couple of useful ways.

First, it is possible to set up a secure "tunnel" within MNET. To set up such a tunnel, two customers select Limited service allowing only the other customer in.

Second, a department can set up an intra-agency server. Their field offices may use "One-way" service and the main office can use "One-way except...," where the except list limits by source address as well as destination address and port.

Keeping it going

Keeping all this information correct is a major task. To help, we have written a number of scripts. These scripts analyze the router configuration and look for inconsistencies, omissions, or additions. Therefore, we can readily catch and correct any mistakes.

Even though MNET is a large network, we are able to work closely with our customers. Because of this closeness, we are aware of who is authorized to request changes (e.g., opening up access lists). If we see a request come in and it is not from the authorized contact, we hold off implementing it until we can verify that it is a valid request.

Further, we require all change requests to be made in writing (e-mail or fax). We retain the original requests and all related correspondence indefinitely. In this way, we maintain an audit trail of how the network was configured and why.

Summary

Our approach to security allows MNET to act as a reliable carrier of data. MNET can certify that the data is not being monitored, that it is correctly delivering data, and that the source of the data can be trusted. Also, the organization responsible for an information resource has the ability to control (at the network level) who can access the resource.

Furthermore, we can perform these functions while avoiding the responsibility of having to certify the security of our customers' internal networks or other connections.

While it is impossible to cover all possible attacks, MNET's approach does allow our customers to perform their operations with a strong sense of security and protects them against most hazards.

References

[1] Garfinkel, Simson & Spafford, Gene, *Practical UNIX and Internet Security*, O'Reilly & Associates, 1996, ISBN 1-56592-148-8.

[2] Holbrook, P., and Reynolds, J., "Site Security Handbook," RFC 1244, July 1991.

[3] Ranum, Marcus, "Internet Firewalls Frequently Asked Questions (FAQ)," Available from: `http://www.iwi.com/pubs/faq.htm`

[4] Ranum, Marcus, "Thinking About Firewalls," 1993. Available from: `ftp://moink.nmsu.edu/firewalls/fwalls.ps.Z`

[5] Doty, T., "The Firewall Heresies," *ConneXions,* Volume 9, No. 6, June 1995.

[6] Doty, T., "A Firewall Overview," *ConneXions,* Volume 9, No. 7, July 1995.

[7] Chapman, D. B., and Zwicky, E. D., "Internet Security Strategies," *ConneXions,* Volume 9, No. 12, December 1995.

[8] Chapman, D. B., and Zwicky, E. D., "Internet Security Policies," *ConneXions,* Volume 10, No. 1, January 1996.

[9] Stallings, W., "Cryptographic Algorithms," Part I: Conventional Cryptography, *ConneXions,* Volume 8, No. 9, September 1994. Part II: Public-Key Encryption and Secure Hash Functions, *ConneXions,* Volume 8, No. 10, October 1994.

[10] Stallings, W., "Pretty Good Privacy," *ConneXions,* Volume 8, No. 12, December 1994.

[11] Kaliski, B., "An Overview of Public-Key Cryptography Standards," *ConneXions,* Volume 6, No. 5, May 1992.

[12] B. Clifford Neuman and Theodore Ts'o, "Kerberos: An Authentication Service for Computer Networks," *IEEE Communications,* Volume 32, No. 9, September 1994.

[13] Levy, Steven, "Clipper Chick," Interview with Dorothy Denning, *WIRED,* 4.09, September 1996, `www.wired.com/4/09/denning`

[14] "IAB and IESG Statement on Cryptographic Technology and the Internet," *ConneXions,* Volume 10, No. 8, August 1996.

[15] Kaufman, Charlie, Perlman, Radia and Speciner, Mike, "The Clipper Proposal," *ConneXions,* Volume 10, No. 9, September 1996.

[16] Kaufman, Charlie, Perlman, Radia and Speciner, Mike, *Network Security: Private Communication in a Public World,* ISBN 0-13-061466-1, Prentice Hall, 1995.

[17] Schiller, J., "Issues in Internet Security," *ConneXions,* Volume 7, No. 9, September 1993.

[18] *ConneXions,* Volume 4, No. 8, August 1990, "Special Issue on Network Management and Network Security."

[19] Bellovin, S.; "Problems in the TCP/IP Protocol Suite," *Computer Communications Review,* Volume 19, No. 2, April 1989.

CRAIG A. FINSETH is the Senior Network Designer for the data network for the State of Minnesota. Previously, he has managed the data networks for the University of Minnesota and the Minnesota Supercomputer Center. Over the past eight years, he has been a major participant in building the Internet within Minnesota. He has been involved in computer networking for nearly twenty years, starting on the ARPANET as a student at M.I.T. He is a member of the Midwest Electronic Crime Investigators' Association and has appeared in court as an expert witness on the Internet. E-mail: `fin@finseth.com`

A Proposed Protocol Improvement for ATM Cell Processing Within SDH Multiplexers

by
Gerard Parr and Stephen Wright,
University of Ulster

Preamble

We are all well aware of the great strides being made in application development (e.g., WWW, 3-D Animation, Virtual Reality, etc.) over the past five years. How many of you have stopped to consider as with the human body the underlying infrastructure that gives "life" to these distributed applications? The infrastructure in this case is that provided by the PTT in a particular country that serves to carry Internet, WWW, voice, telephony and ATM sessions around the developed and developing world. That infrastructure is better known as SDH a combined world-wide multiplexing structure capable of gigabits per second throughput. Maturing as with all "maturing standards" there are aspects of their definition that evolve, sometimes by chance, but usually out of practice in the field. The purpose of this article is to inform and promote awareness amongst the networking community of recent research under the Telecommunications and Distributed Systems Research Group (T&DSRG) of the University of Ulster that aims to improve the performance of SDH multiplexers as they access and decode the ATM cell payload of STM-1 frames. This work has International applicability in the Telecommunications Industry and should be of interest to all those who develop and manufacture underlying transmission and switching technology for BISDN.

Introduction

In a SONET/SDH environment ATM cells exiting a switch enter the tributaries of a local multiplexer which maps incoming cells into containers to be transported in a frame structure within a high speed link (155Mbps, 622Mbps, or 2.5Gbps) to the next upstream or downstream network multiplexer. On arrival to this multiplexer, the cell payload is extracted and passed to the switch on the multiplexer tributary side where local cells are switched to their destinations and non-local cells are switched to the multiplexer where the aforementioned operations are performed. These functions are performed at each intermediate multiplexer and switch until cells arrive at their final destinations. This would imply that unnecessary processing of cells is occurring at multiplexers and switches and is therefore causing a burden to both devices. Cell processing could be improved if for example it was known that all cells within a container(s) were destined for the next upstream/downstream node and could be forwarded without extraction. This article focuses on ATM cell processing at both the network switch and multiplexer. The present method of multiplexing within the Synchronous Digital Hierarchy (SDH) will be discussed and a proposed improved schema will be introduced.

Multiplexing

There are two defined multiplexing standards, one the *North American* (NA) standard, used mainly in the USA, Canada and Japan. The other, the *European* standard covers the rest of the world. The NA standard uses T1 circuits at 1.54Mbps (24 × 64Kbps channels), and the European standard uses E1 circuits at 2.048Mbps (32 × 64Kbps channels). Digital systems must convert voice traffic into digital code by periodically sampling the voice traffic, a process known as *Pulse Code Modulation* (PCM) encoding. Both standards use the same sampling rate, i.e., each incoming channel is sampled at a rate of 8,000 times per second, one byte (8 bits) every 125 microseconds (ms), however, each use incompatible devices making interworking expensive.

The process of combining multiple channels together is called *Time Division Multiplexing* (TDM) [1]. Each sampled byte is multiplexed to a high-speed link, and the position occupied by each byte in the link is called a *timeslot*. The first level of multiplexing (primary) for an E1 link multiplexes thirty incoming 64Kbps voice channels. At the receiving end of the E1 link, the de-multiplexer must have the capability of recognising each byte within each channel. These 30 channels constitute a 32 byte frame. The additional 2 bytes, one of which identifies the beginning of the frame, is called the *Frame Alignment Word* (FAW), the other contains signalling information for up to 30 telephone circuits. Timeslot 0 (TS0) holds the FAW, TS1–TS15 hold data sampled from channels Ch1–Ch15. TS16 holds the signalling information, and sampled data from channels Ch16–Ch30 are held in TS17–TS31. The contents of a frame may consist of 30 voice channels, or 31 data channels. Identification of a particular channel is done by identifying the FAW, then counting to the appropriate channel. The NA primary frame can only multiplex 24 channels and does not contain any FAW. The contents of a NA frame may be either 24 × 64Kbps voice channels or 24 × 56Kbps data channels.

The Plesiochronous Digital Hierarchy

As a result of the growth in demand for voice traffic within the *Public Switched Telephone Network* (PSTN), the *Plesiochronous Digital Hierarchy* (PDH) evolved. Initial multiplexing levels both in Europe and the US were becoming inadequate due to increasing traffic levels. Further levels of multiplexing were introduced to both standards, in order to reduce the number of primary level (First Order) multiplexers used (1.54Mbps, and 2.048Mbps). In Europe, the next level of multiplexing was 4 × 2.048 giving a 8.448Mbps link (Second Order), then 4 × 8.448Mbps giving 34.368Mbps (Third Order), 4 × 34.368Mbps giving 139.264Mbps (Fourth Order). Multiplexing at these levels has its problems. In theory, incoming channels for multiplexing from primary multiplexers run at exactly the same speed, but, this is not the case as some signals may be generated slightly faster or slower than others, depending on the generating synchronising clock speeds. These rogue signals therefore need aligning to the same bit rates by bit stuffing justification bits, a process known as *plesiochronous* (almost synchronised) operation. When the signal is received by the de-multiplexer it readily recognises these bits and discards them. A national clock controls the synchronisation of equipment in the network by using a master clock with a varying frequency of no more than 1:1011 as specified by CCITT/ITU-T. The network nodes extract this signal from the clock via dual timing paths and forward to any synchronous equipment they may control. If both these timing paths should fail, the nodes will use their internal (slave) clocks to provide suitable timing signals. All higher rates of multiplexing above 1.54Mbps and 2.048Mbps use bit interleaving which contributes to a multiplex mountain. This problem is prominent when providing a new customer connection to a 2.048Mps leased line (drop and insert). This scenario means that a higher rate e.g., a 140Mbps link has to be de-multiplexed down to 2Mbps to connect the customer, then multiplexed back up again to 140Mbps [2]. This process leads to problems of inflexibility in connection, provision of service and control due to the amount of equipment used. The two main deficiencies with the PDH are the inability to recognise individual channels in a high-speed bit stream, and the frame structure used has insufficient provision for carrying network management information.

continued on next page

Improving ATM Cell Processing *(continued)*

These problems may not be so important within the PSTN, but with the introduction of new services such as LAN interconnection, videoconferencing, image database browsing, real time video, etc, the need for better control and management is crucial.

The Synchronous Digital Hierarchy

A recent multiplexing standard, the *Synchronous Digital Hierarchy* (SDH), has evolved from the *American National Standards Institute* (ANSI) *Synchronous Optical Network* (SONET) standard [3–5]. The main CCITT/ITU-T recommendation covering SDH is G.707 which defines the bit rates, covers both the Network Node interconnection issues and multiplexing structure. This single recommendation replaces the previously defined (G.707, G.708 and G.709). Working groups have been established to cover other issues such as optical interfaces, and *Operations Administration and Management* (OAM). The fundamental transmission rate within SDH is 155.52Mbps, defined as a *Synchronous Transport Module 1* (STM-1). Higher rates are four times and sixteen times the fundamental giving 622.088Mbps (STM-4), and 2.48Gbps (STM-16), with further higher levels under study. STM-1 signals are defined to carry lower signal rates, allowing the existing PDH signals [6], (with the exception of 8Mbps, which is still under study), to be carried over this synchronous network as STM-1 frames, a major reason why SDH was readily accepted. Access to the SDH is via a network node, or *network element* (NE) [7]. A NE is any device capable of recognising and maintaining a 155Mbps signal, e.g., an STM-1 multiplexer. An STM-1 terminal multiplexer is capable of multiplexing 63 × 2Mbps signals every STM-1 frame, it is of modular construction, with its main processing carried out within its core, which contains the *Payload Manager* (PM) which functions include placing *Virtual Containers* (VCs) into STM-1 frames and to calculate any necessary pointers. Outgoing signals (aggregate) on the STM-1 link side are scrambled, and undergo an electrical to optical conversion if necessary. This port is capable of shutting down in the event of an optical signal loss. A synchronous multiplexer performs both the multiplex and de-multiplex functions for signals and for this reason it is referred to commonly as a MUX. Within the same location of the multiplexer(s), a local *Element Manager* can connect to individual multiplexers via a LAN using a CCITT defined Q interface [2]. Remote element managers can communicate to distant multiplexers by way of the *Embedded Communication Channel* (ECC) within an STM-1 frame using reserved MSOH octets, or via an X.25 channel connected to the manager using a Q interface.

STM-1 frame structure

An STM-1 frame structure, the basic unit of the SDH, consists of a single *Administrative Unit Group* (AUG), and a *Section Overhead* (SOH), and has a duration of 125ms, with a total contents of 2,430 8-bit bytes (octets). The STM-*N* frame consists of $N \times$ AUGs plus a SOH. Eighty-one octets are reserved for the SOH, which remains in the frame between two adjacent synchronous multiplexers, allowing the allocation of channels for such functions as OAM, user channels, protection switching, section performance, and frame alignment. The remaining 2,349 octets are for traffic (payload). Payload within an STM-1 frame repeats every 261 octets. An STM-1 frame may be viewed as a 9 row by 270 column matrix, which can be read from top to bottom and left to right. At the beginning of the SOH at row 1, there are 6 octets reserved for the FAW. Payload can be inserted into or dropped from the STM-1 at any node equipped with a drop and insert multiplexer. In the event of a phase variation in the signal, the payload may shift, and as a result, it is allowed to float within the frame, thus the need for pointers.

The higher rates within the SDH are achieved by byte interleaving, thus overcoming the limitations of bit interleaving within the PDH, providing easy access to a new customer, without the need for a multiplex mountain, allowing more efficient drop and insertion of channels, leading to the provision of higher bandwidths for multimedia services and easy access, lower operating costs, and less multiplexing equipment.

SDH and SONET

Although similar, these systems use different transmission rates, terminology, and hierarchies. The fundamental transmission rate within SONET is 51Mbps, whereas SDH uses 155.52Mbps. Frames within SONET are termed *Synchronous Transport Signals* (STS-s). The STS-1 structure uses 9 row octets × 90 octets, with the first 3 octet columns used as the section and line overhead. The remaining 87 columns are used to carry the *Synchronous Payload Envelope* (SPE). Payload consists of a DS-3 signal, or lower order signals such as 1.5Mbps, or 6Mbps. The STS-*n* overhead contains a pointer to indicate the position of the SPE. The combination of the payload and the pointer is called an *AU.* A 45Mbps signal is carried within an AU-3. When the STS-1 is converted and scrambled, an *Optical Carrier 1* (OC-1) results, and is used to carry the signal over fibre. A *Virtual Tributary* (VT-*n*), analogous to an STM-1 TU-*n* has four sizes: VT-1.5 which carries a 1.54Mbps signal, VT-2 which carries a 2Mbps signal, a VT-3 which carries a 3.15Mbps signal (48 voice channels), and a VT-6 which carries a 6Mbps signal. Path overheads are included within the VT-*n*s. Both the OC-*n* and STS-*n* signals have been defined separately by ANSI. Three multiplexed STS-1 signals contain 3 × 90 i.e., 270 columns, the format of the STM-1 frame. Three SPE within the STS-3 when concatenated can be regarded as a single entity called a STS-3c, enabling the signal to be transported intact across either SONET or SDH. SONET systems can carry 3×34Mbps signals via an STS-3c signal, and an SDH system can carry 3×AU-3 signals within an STM-1. SONET is therefore viewed as a subset of SDH.

Management of incoming signals to SDH Frames

Incoming signals are mapped into containers, with each container type specifying an existing PDH signalling rate. Depending on the phase variation of the incoming signal, justification bits may be added. Control information is added to each container, a *Path Overhead* (POH), the contents of which may be used to identify and locate a specific signal within an STM-1 frame and helps in monitoring signal performance. Also, containers containing, e.g., only stuffed payload could be identified by the receiving multiplexer, and ignored, thus alleviating some processing overheads. The POH exists between two ends of a synchronous path and provides end-to- end management functionality. A pointer is then added to indicate the position of the VC-*n* within the STM-1 frame.

Multiplexing structure for 2Mbps

A 2Mbps asynchronous signal is mapped into a C-12 and a POH is added to form a VC-12. Additional fixed stuff bits and bytes maintain a fixed size of 140 bytes for a 500ms TU multiframe, i.e., 4×STM-1 frames. Asynchronous mapping allows for justification of the tributary, allowing for variations between the tributary clock rates and the network's synchronous clock. An ATM cell stream may be mapped to a C-4 with its octet boundaries aligned with the C-4 octet boundaries. The C-4 is then mapped to a VC-4 together with a VC-4 POH. As the capacity of the C-4 is 2340 octets, it is not an integer multiple of the ATM 53 octet cell, (44.15 cells), a cell may cross the C-4 boundary.

continued on next page

Improving ATM Cell Processing *(continued)*

Prior to mapping, the ATM cell payload field is scrambled, and when extracted it is unscrambled before passing to the ATM layer. Scrambling is necessary in order to provide security against false cell delineation and payload field replicating the STM-*N* frame alignment word.

Network scenario and operation

Take an example scenario (see Figure 1), where remote LANs e.g., Ethernet are communicating via an ATM network using the Synchronous Digital Hierarchy (SDH) as its physical layer. Each subnet within the proposed network scenario contains 4 CSMA/CD nodes, each generating a variable rate of traffic to every other subnet via an ATM switching network. Packet traffic from these nodes enter a *Interworking Unit* (IWU) which contains an addresses table for its local nodes. A packet whose address is not contained within the table is forwarded to an *ATM Access Unit* (AAU). Functions of the AAU consist of convergence (protocol conversion), packet segmentation/reassembly, and congestion control. Packets now in ATM cell format are passed to a switch which routes them to their *Virtual Channels* (VCs), within appropriate *Virtual Paths* (VPs). Cells now enter a *synchronous multiplexer* (SMUX) which functions will be described shortly. Cells destined for the network enter the tributaries of a multiplexer, are mapped into containers, have various aligning and multiplexing functions applied, a *Path Overhead* (POH), and *Section Overhead* (SOH) added before eventually forming an STM-1 frame, the basic unit of the SDH. STM-1 frames are then transported to the next upstream/downstream multiplexer where the SOH, POH are removed, cell payload extracted and sent to the local switch on the tributary side where cells are either dispatched to their final destinations, or routed back to the same multiplexer for mapping into containers etc., before forwarding to the next multiplexer.

On closer examination, this scenario would imply that there is unnecessary cell processing at multiplexers in the case where a full payload is destined for the same end destination node that is not accessible via the current multiplexer tributaries. If multiplexers were provided with additional "intelligence" enabling them to identify such payloads then this operation would avoid unnecessary payload extraction and in the same instance alleviate cell processing at switches.

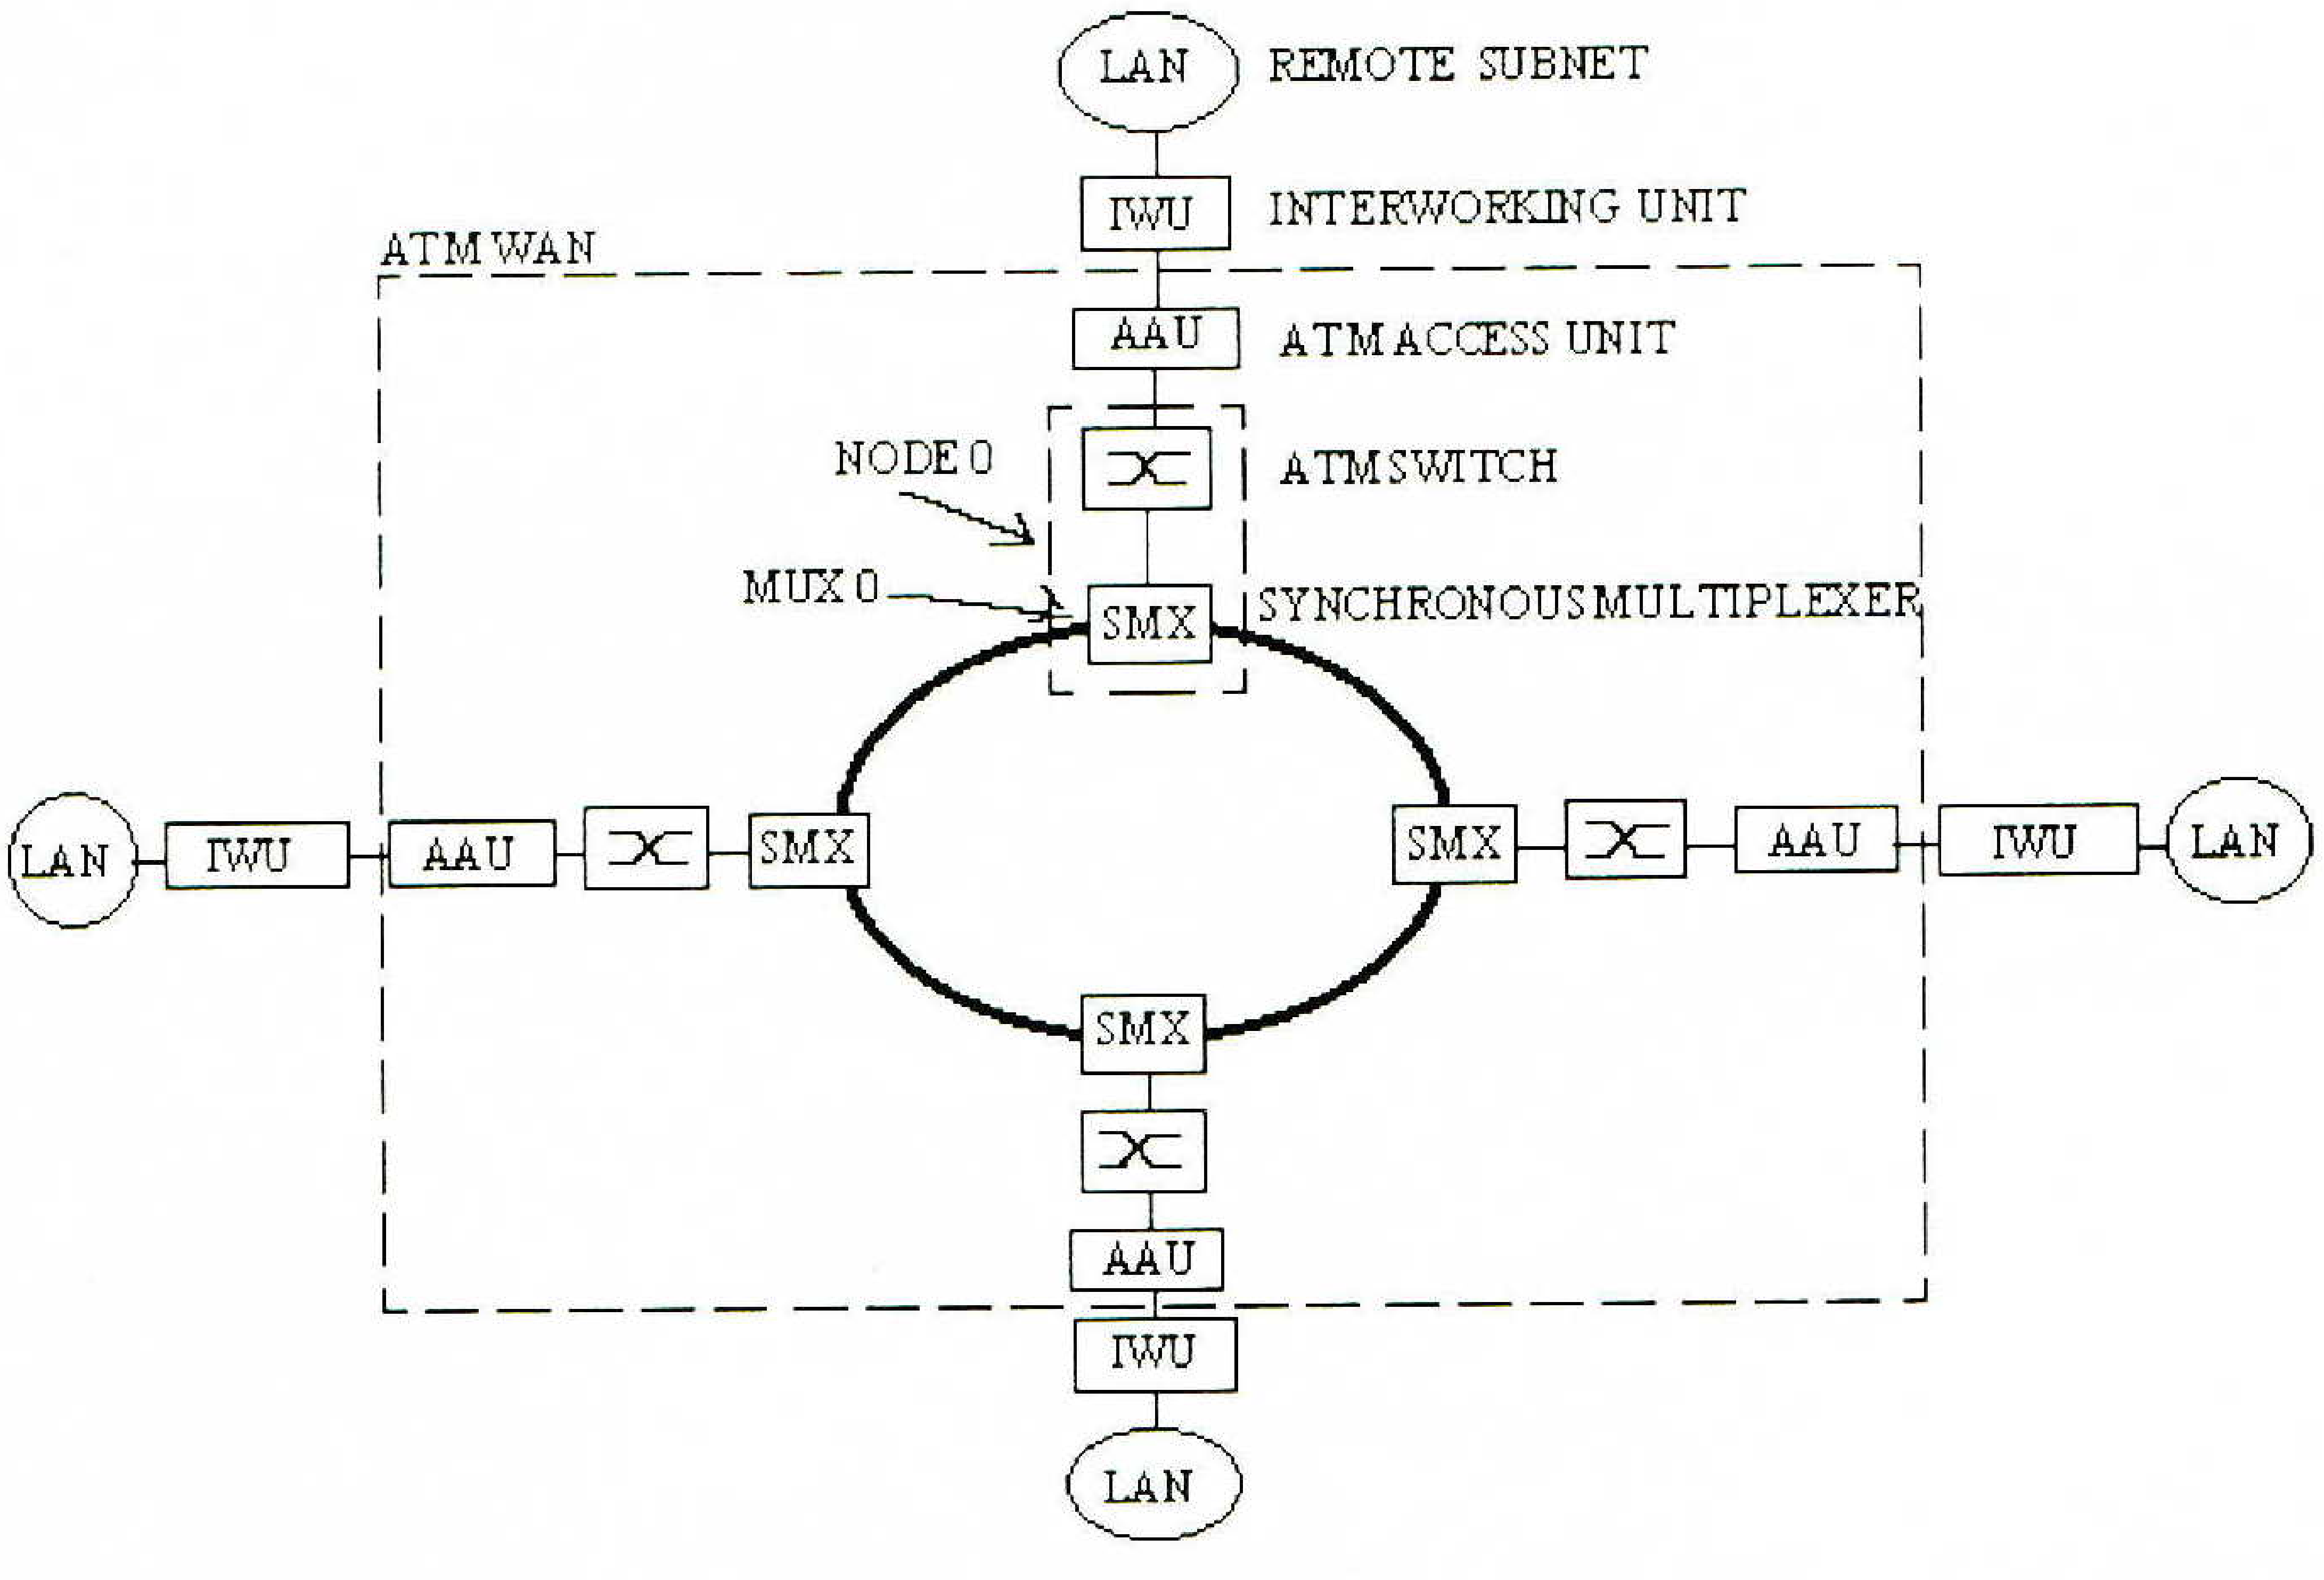

Figure 1: Proposed Network Scenario

Current multiplex Operation for ATM cells

CCITT/ITU-T recommendation G.707 defines the SDH multiplexing structure and allows the ATM cell stream to be mapped into C-4 containers. As the payload contents of the C-4 is not an integer multiple of the cell size of 53 octets, each C-4 holds 44.15 cells but allows for cells to overlap a C-4 boundary. The use of other container types for cell mapping is presently under study. A POH is added to the C-4 to form an VC-4, and is used to trace and monitor the payload throughout its journey. A SOH is now added to form a STM-1 frame. The SOH terminates at every intermediate MUX as the information it contains relates to the section between two multiplexers. STM-1 frames arriving at a MUX have their payload extracted and passed to the tributary side of the multiplexer where they enter a switch which may pass the cells to its local AAU for onward transmission to their final destination or route them to the multiplexer for mapping into C-4s, etc. This operation repeats at every intermediate MUX until cells are delivered to their final destinations.

Weaknesses with current multiplex operations

From the previous mentioned multiplexing/de-multiplexing operations, it can be seen that unnecessary and time consuming processing overheads occur at intermediate MUX nodes with the extraction of payloads, which may be avoided if greater "intelligence" could be added to multiplexers, enabling them to detect payloads which are the same and not destined for any of the local tributaries. This being the case, a new SOH could be generated and the intact container forwarded to the next upstream/downstream MUX thus alleviating processing burdens on both the MUX and local switch.

Proposed MUX operations

Within the VC-4 POH there are three octets, Z3, Z4 and Z5 which (at time of writing) are either partially defined or remain undefined. The Z3 octet is used for user communication purposes as is the F2 octet. The Z4 octet is undefined, and the Z5 octet is partially defined for management purposes e.g., error count. What is proposed is that these octets should be utilised as follows: Cells arriving at the tributary side of the multiplexer destined for a particular VCI, VPI, are mapped to a container of cells. Given that the multiplexer would have captured previous VCI, VPI setup data and stored it in a state information table, two header fields of every incoming cell i.e., their VCI,VPI may be examined and from this information the multiplexer can determine if the contents of the current container is for the same destination. If it is, the Z5 octet of the POH is set and the VCI, VPI values are stored within the Z3, Z4, octets indicating to multiplexers upstream/downstream that this entire payload is for the same end device.

Subsequently, when a remote multiplexer receives an STM-1 frame and strips the SOH, the Z5 octet of the POH is examined, and if it is set, the Z3, Z4 octets are then examined respectively for the destination VCI, VPI pair, and these values are then compared with the state information table data. If the pair match for a local tributary circuit, then the payload is extracted and the cells forwarded on the relevant tributary. Otherwise a new SOH is added and the entire payload within the STM-1 frame is transmitted to the next multiplexer.

Implementation of proposed algorithm

As previously mentioned, on the identification of a full container the Z3, Z4, and Z5 octets of the VC-4 POH could be used to determine the handling of a C-4 payload, i.e., extract or leave intact. The Z5 octet denoting that the container load is destined for the same end point, and the Z3 and Z4 octets representing the associated VC and VP pair respectively. Within the model, it is only necessary to use two of the proposed octets, i.e., Z5 to denote that the payload is destined for the same end point, and Z4 denoting the end point address.

continued on next page

Improving ATM Cell Processing *(continued)*

On the receipt of a STM-1 frame at a multiplexer, the SOH is stripped, the POH Z5 octet is selected and examined, and if set to 1, the Z4 octet is compared with this device's identity value. If recognised as this multiplexer's address then the payload is extracted and the cells are passed to the switch for routing to the local IWU, otherwise a new SOH is generated and the STM-1 frame is forwarded to the next upstream/downstream multiplexer. The modules constructed to implement the proposed algorithm will now be explained.

Proposed operation of C-4 and POH modules

As each incoming cell enters the module, a *Test Memory* module is read to see if any other cells have entered in this container cycle. If this is the first cell in, then the *Test Memory* module is set to 1, the destination node is searched for within the cell routing vector and the value is written to a memory module *Dest Mem*. Further incoming cells have their destination node address tested with the current value in the *Dest Mem* module, if they differ then the memory module *AllSame Mem* is reset to 0, thus indicating that this container has a mixed payload and is not destined for the same end node. All incoming cells are mapped into the C-4 as per the conventional method, including the generation of empty cells. When a container is full, (i.e., 44 cells, this must be an integer value for modelling), the *AllSame Mem* module is read and if the value is 1, the Z5 octet is set, the Dest Mem value is placed into the Z4 octet (i.e., the identity of the destination node), and the VC-4 is passed on as before. It should be noted that in the real world, the VCI/VPI pair would be stored in a table and then written to the Z3, Z4 octets as stated.

Proposed operation of STM-1 frame reception

Received STM-1 frames have their VC-4 extracted and the Z5 octet is examined. If set, the Z4 octet is read, and if the value matches this node's identity, the cells are extracted. What has to be remembered in this proposed scenario is that if these cells have remained within their container for more than 1 hop then the routing vector will not have been updated. In order to check this phenomenon the cell Route Vector must be adjusted before passing to the switch.

Simulation model

A model containing two proprietary ATM switch architectures operating over an SDH physical layer was built, in accordance with existing CCITT/ITU-T standards. The modelled WAN representation can be seen in Figure 1. When cells arrive at the trib side of the multiplexer destined for a particular VCI, VPI, it creates a container and clocks in 44 cells. Given that the multiplexer would have captured previous VCI, VPI setup data and stored it in a state information table, it now can examine two header fields of every incoming cell (i.e., their VCI,VPI) and determine if the batch of 44 cells are for the same destination. If they are, it sets the Z5 byte of the POH and stores the VCI, VPI values within the Z3, Z4, indicating to multiplexers downstream that this entire payload is for the same end device. Subsequently, when a remote multiplexer receives an STM-1 frame and strips the SOH, it then examines the Z5 byte of the POH. If it is set, it examines the Z3, Z4 bytes respectively for the destination VCI, VPI pair and compares these values with the state information table. If the pair match for a local trib circuit, then it disassembles the payload and forwards the cells on the relevant trib. Otherwise it adds a new SOH and transmits the entire payload within the STM-1 frame to the next multiplexer. At simulation start time, traffic is generated at a random rate from the LANs and enters an Interworking Unit (IWU) module. Packets destined for the ATM network are passed through the IWU to the network nodes.

The network nodes consist of a multiplexer, switch, and an ATM Access Unit (AAU) with functions including convergence, segmentation, and cell sequencing. Segmented cells, carry their own routing (path) information, which is computed within the segmentation module. Cells received at switches undergo a mapping process related to their next hop and are switched to their appropriate virtual path (VP). Cells entering the multiplexer are now multiplexed, to form an STM-1 frame, consisting of [C-4 container, Virtual Container (VC), Administrative Unit (AU), and Section Overhead (SOH)]. The transmission delay is modelled by delaying the frame within the link. Simulation parameters were initialised in order to change the load on the network.

Simulation results

Experiments were conducted to ascertain the difference in performance between the conventional method of multiplexing and the proposed method in order to observe:

- Cells generated from any source were received at the correct destination.
- Mean end-to-end delay of cells.
- Number of cells passing through the MUX both dummy and live.
- Hop count of cells.
- Number of VC-4s having their Z5 octet set.

The latter i.e., the number of VC-4s having their Z5 octet set only applies to the new proposed method simulation. The overall number of received cells at the subnets was found to equal the number of generated cells. In reference to Figures 2 and 3, it can be observed that there is a marked decrease in the mean end-to-end cell delay with the new approach. It may be observed from both these figures that mean end-to-end cell delay increases as the simulation time increases, this is due to the increase in traffic intensity i.e., load, which should be the case.

A comparison of both figures shows that an improvement in end-to-end cell processing has been achieved. In keeping with the requirement to maintain a stable network, as the payload would have been extracted using the conventional method, a complete payload will need to be generated for the next outgoing container. In the event that cells do not arrive within a time-out period from the trib side of the MUX then dummy cells will have to be generated and mapped into the container's payload until the arrival of live cells or the container is full. Under the new method, the need for such cell stuffing is reduced, or in some cases eliminated i.e., where the POH Z5 octet is set. The results show that with the proposed method a reduction in dummy cell generation had been noticed, an indication that a reduction of unnecessary cell processing had been achieved. Cells passing into a switch from the MUX will have their hop count field incremented. If a cell coming from a switch to the MUX has a hop count greater than zero, i.e., the cell has entered the switch and was found not to be for a node on the switch side and it is switched to the network via the MUX, it will have a hop count of at least 1. This count is recorded within a module in the MUX. The number of cells with hop count greater than zero was recorded and from these figures it was observed that the number of cells with hop counts greater than zero was less for the proposed method, which indicates that the amount of cell processing has again been reduced.

continued on next page

Improving ATM Cell Processing *(continued)*

The results overall show a reduction in the number of hop-counts for cells and show on average an overall processing gain of 35%, an indication that the proposed algorithm may be applied to the functions of the ATM/SDH environment to alleviate unnecessary cell processing at both ATM switches and SDH multiplexers.

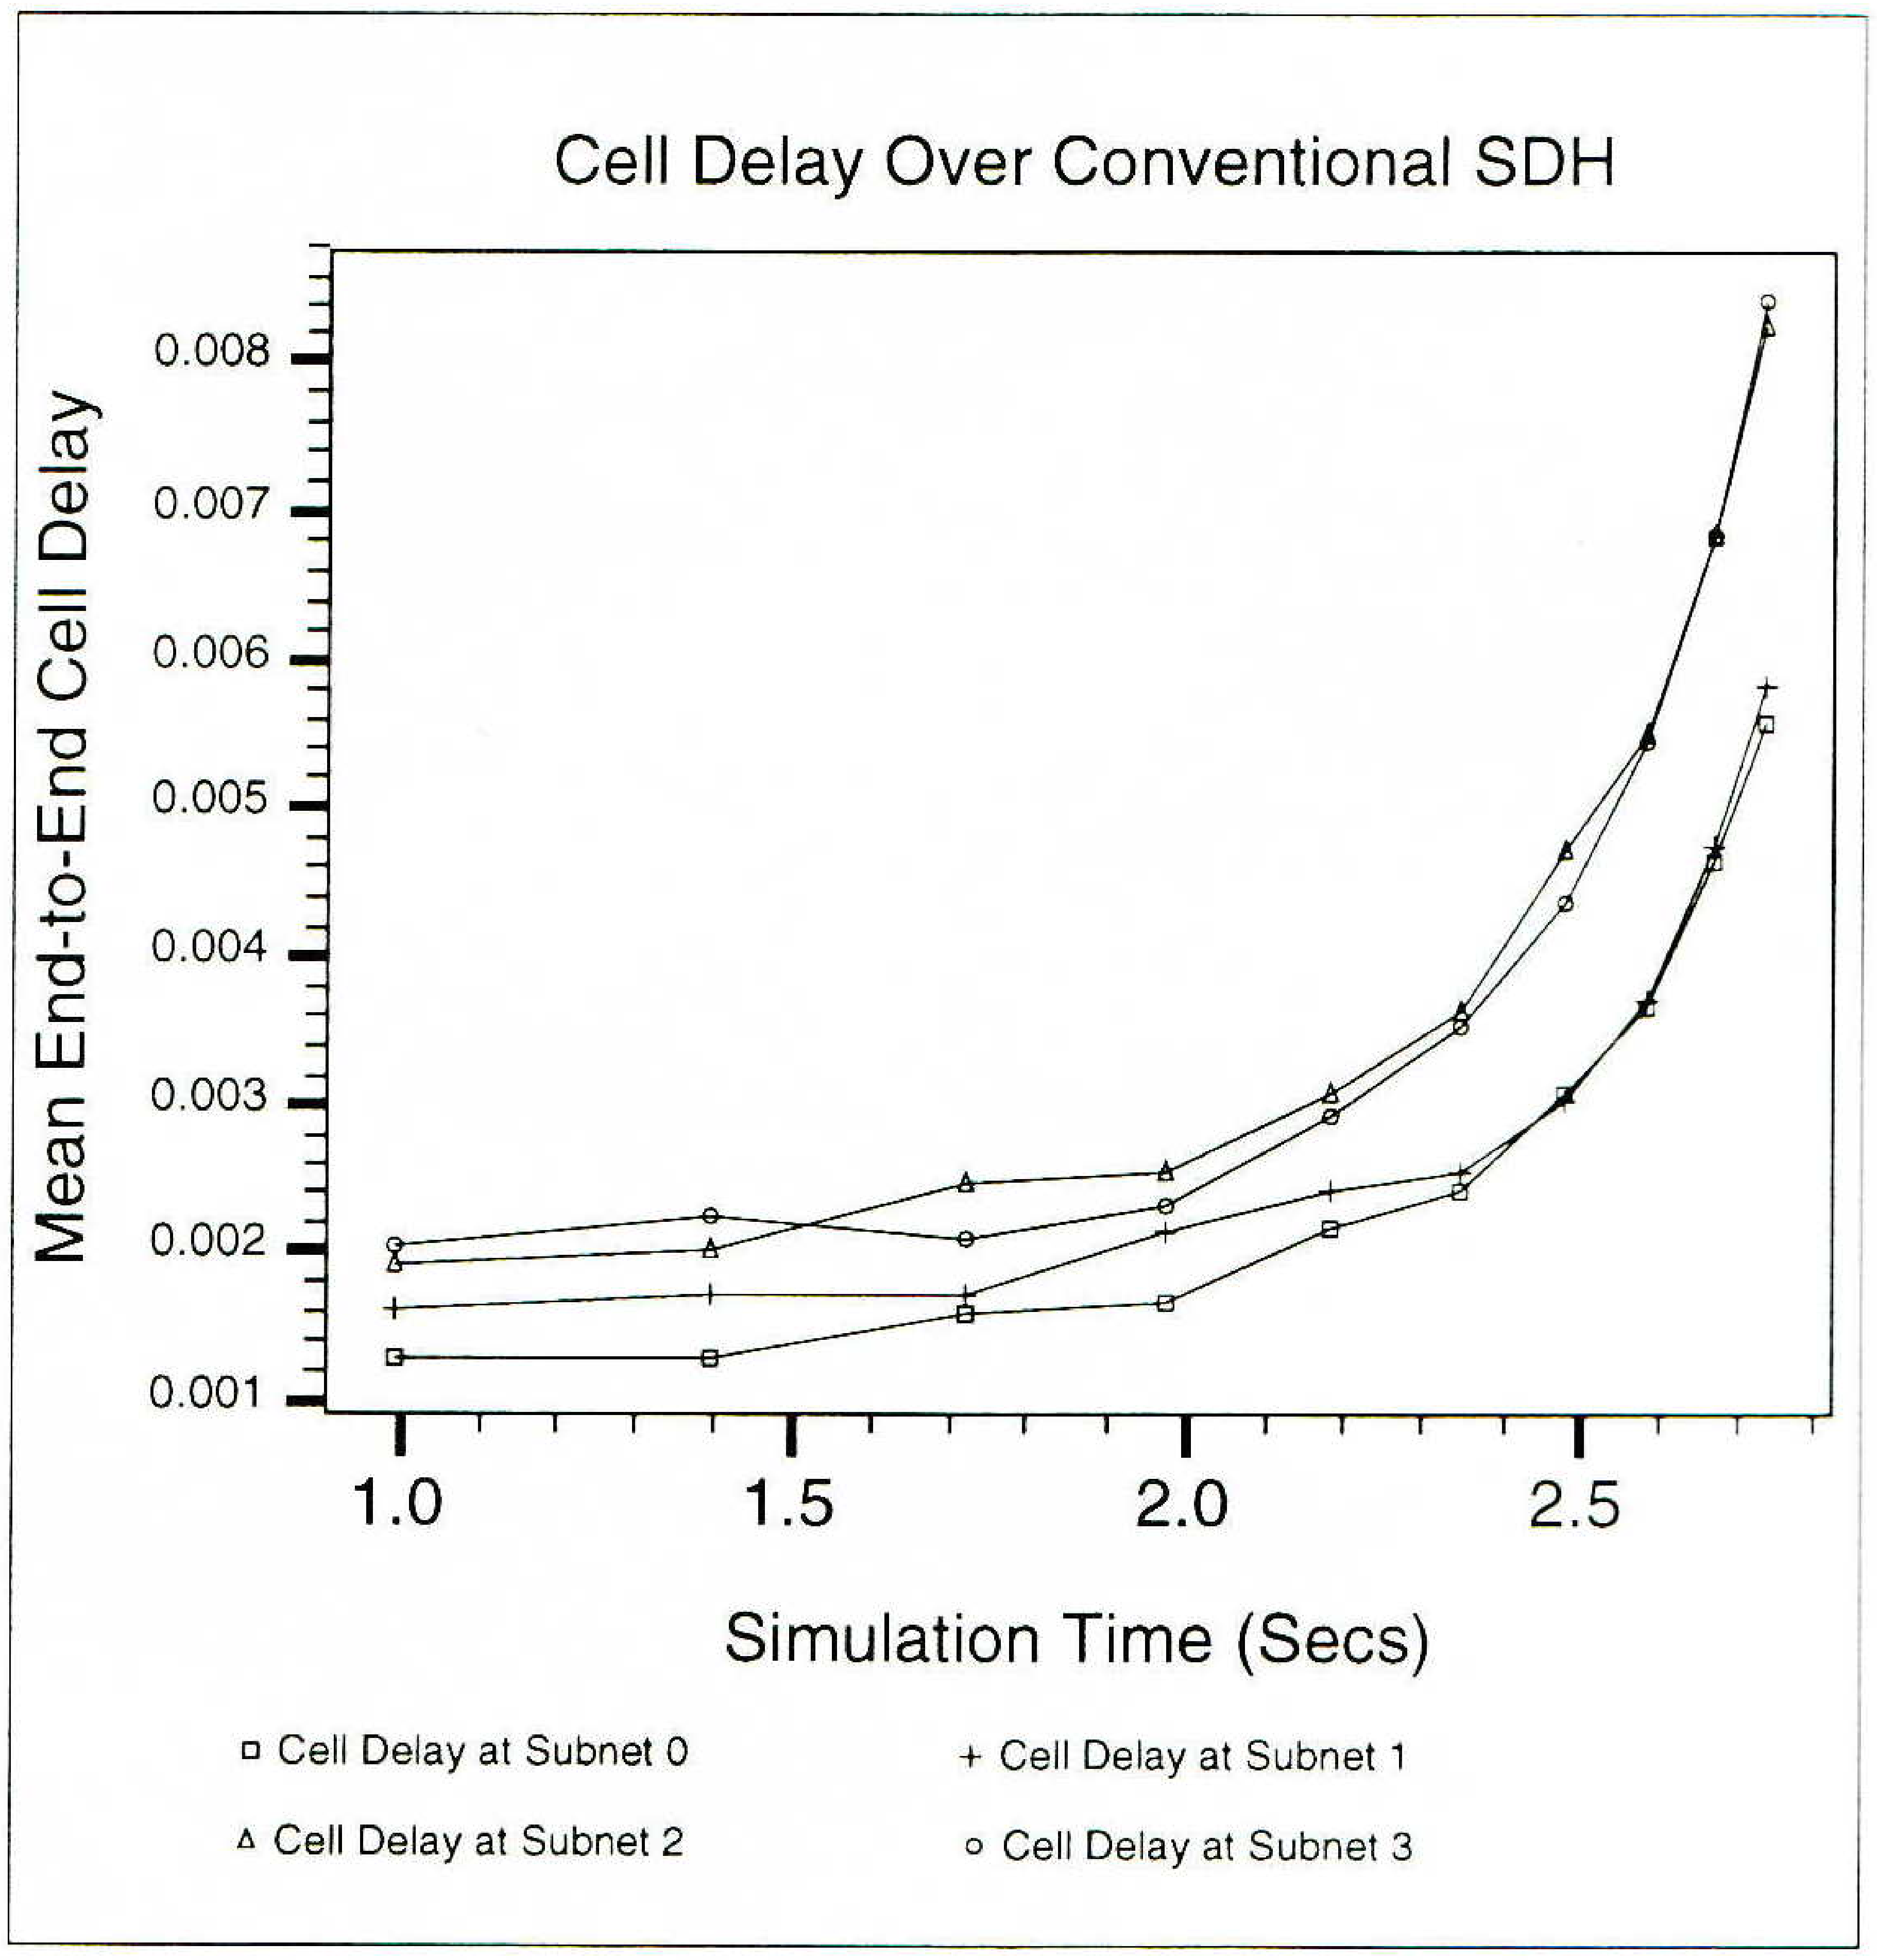

Figure 2: Mean End-to-End Cell Delay (Conventional Method)

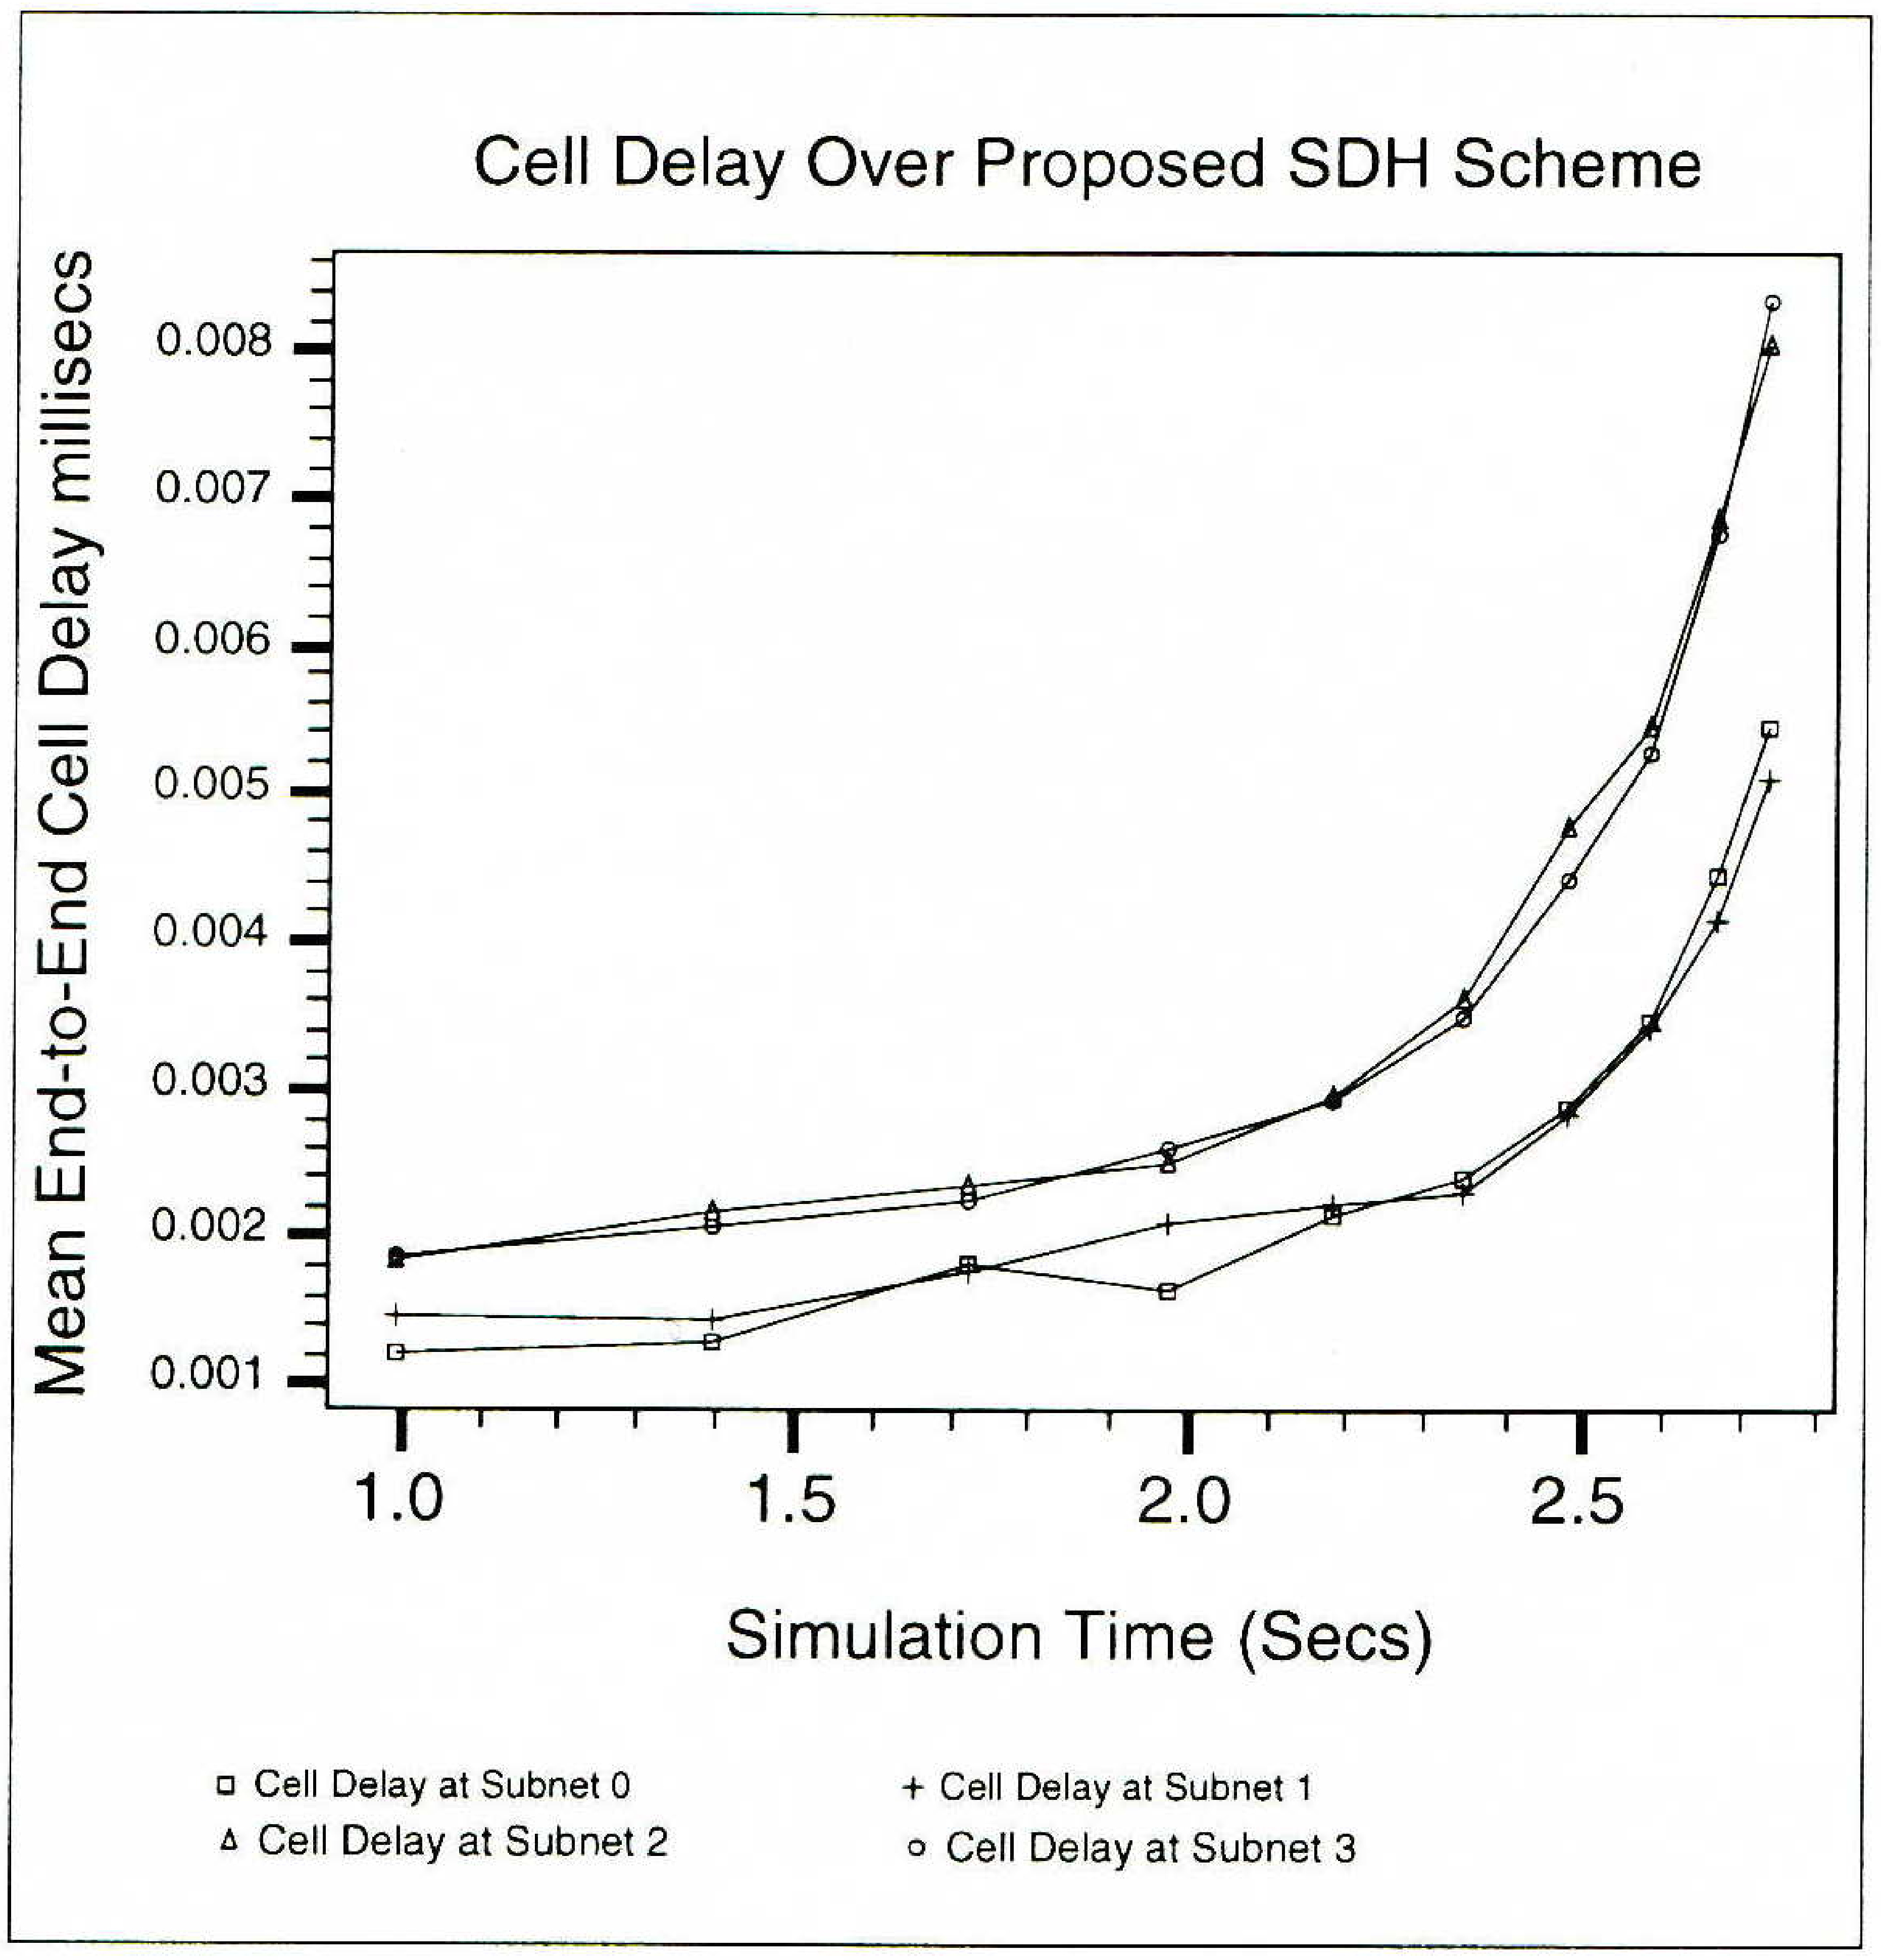

Figure 3: Mean End-to-End Cell Delay (Proposed Method)

Conclusions

To overcome the problems associated with undue processing overheads at ATM network switches and SDH multiplexers, this article proposed a simple schema for using the partially defined Z3-Z5 bytes of the STM-1 frame VC-4 POH. An overview of the modelled WAN was given and a description of experiments and results were discussed for both the conventional and proposed methods of multiplexing. Data were gathered for the following reasons:

- To ensure source generated cells arrived at their destinations.
- The mean end-to-end cell delay.
- A count of dummy and live cells passing through the MUX.
- A hop count for cells.
- The number of VC-4s having their Z5 octet set.

The above criteria was used to test and compare both conventional and proposed methods in order to show that with the use of the proposed method, an improvement could be obtained in cell processing at both the MUX and the switch. From the results obtained, an improvement was seen in all gathered data thus showing that the proposed algorithm could be applied to the functions of the ATM/SDH environment and alleviate the burden of unnecessary cell processing within both devices. In the wider context, this work has far reaching implications for PTTs and manufacturers of multiplexers. Moreover, the "cost" of implementing this new protocol in future MUX designs is minimal in software engineering terms, whereas the advantages in larger systems is significant. These advantages can be summarised as follows:

- Reducing unnecessary loading on ATM switches;
- Reduction in the risk of cell loss within switches;
- More efficient use of SDH infrastructure bandwidth (e.g., less dummy cell traffic);
- Reduction in processing overheads within SDH multiplexers at the network connection;
- More efficient delivery of payload, which in real terms is the user application traffic.

While the work presented here is based around the "physical" engineering transport layer it seems to demonstrate the tight coupling that exists between underlying PTT equipment performance and end-user application support.

References

[1] Halsall, Fred, *Data Communications, Computer Networks And Open Systems,* 3rd Edition, Addison-Wesley., 1992.

[2] Mathews M., et al, "Synchronous Transmission Systems," Northern Telecom Europe Limited, 1993.

[3] Ballart Ralph, Ching Yau-Chau, "SONET: Now It's the Standard Optical Network," *IEEE Communications Magazine,* March 1989, pp 8–15.

[4] Boehm Rodney J., "Progress in Standardization of SONET," *IEEE LCS—The Magazine of Lightwave Communication Systems,* May 1990, pp 8–16.

[5] Hac Anna, Matlu Hasan B., "Synchronous Optical Network and Broadband ISDN Protocols," *Computer,* November 1989, pp 26–34.

continued on next page

Improving ATM Cell Processing *(continued)*

[6] ITU-T Recommendation G.702, "Digital Hierarchy Bit Rates," 1993.

[7] ITU-T Recommendation G.784, "Synchronous Digital Hierarchy (SDH) Management," 1994.

[8] Parr G., Wright S., "Reducing ATM Cell Processing in SDH Multiplexers," 1996 IEEE AFRICON, University of Stellenbosch, South Africa, September 24th–27th, pp 118–123.

[9] Wright S., "An Algorithm to Improve ATM Cell Processing in SDH Multiplexers," D.Phil. Thesis, September 1995.

[10] Laubach, Mark, "To Foster Residential Area BroadBand Internet Technology: IP Datagrams Keep Going, and Going, and Going..." *ConneXions,* Volume 10, No. 2, February 1996.

[11] Laubach, Mark, "ATM for your internet—But When?" *ConneXions,* Volume 7, No. 9, September 1993.

[12] Atkinson, Ran, "Towards Real ATM Interoperability," *ConneXions,* Volume 7, No. 8, August 1993.

[13] Laubach, M., "IP Over ATM and the Construction of High-Speed Subnet Backbones," *ConneXions,* Volume 8, No. 7, July 1994.

[14] Clapp, G. & Zeug, M., "Components of OSI: Asynchronous Transfer Mode," *ConneXions,* Volume 6, No. 4, April 1992.

[15] Rybczynski, Tony, "The ATM Enterprise Network Switch: Enabler of Network Transformation," *ConneXions,* Volume 9, No. 11, November 1995.

GERARD PARR graduated from the University of Ulster with a BSc (Hons) in Computer Science, and a PhD in Telecommunications & Computer Networking in 1984 and 1990 respectively. He is actively involved in various European (e.g., STRIDE, Telematique, Telematics, INTERREG II) and locally-funded projects (e.g., BT University Development Awards), both within the University, and in conjunction with large industrial partners. The subject matter for these projects ranges from state of the art communications research into Generic ATM Switch Fabric Designs to support Broadband-ISDN application environments, Bridge Protocol Development, SDH Network Management, Distributed Geographical Information Systems and the quest for the Common ISDN-API. He has published widely in Ireland, the UK, and the USA in numerous journals including BCS, IEEE and ACM. He is also an advocate for technology transfer, and in promoting the strategic nature of Telecommunications in support of BioInformatics, Lingumatics and inward investment. E-mail: `gparr@causeway.infc.ulst.ac.uk`

STEPHEN WRIGHT is currently a research officer in the Telecommunications & Distributed Systems Research Team at the University of Ulster (Coleraine). He received a BSc. (Hons) in Computer Science in 1991 and a PhD. in Telecommunications Engineering in 1996 from the University of Ulster. He is currently involved in NIDevR funded research, examining ATM Services and Support Mechanisms. He has specific interest in SDH Network Modelling and ATM Switch Fabric design. Dr. Wright has published in BCS and IEEE.
E-mail: `swright@causeway.infc.ulst.ac.uk`

This research has been funded by the EU STRIDE programme and by the Northern Ireland Development & Research Fund within the Department of Education, Northern Ireland.

100VG and 100-BaseT Tutorial

by Roger Cohen, Netmarq

Introduction

100VG-AnyLAN and *100-BaseT* are two different local area network (LAN) technologies that operate at a rate of 100 Megabits per second (Mbps). They are competing to become the de facto standard for new LANs and for upgrades to existing Ethernet LANs and represent two contrasting solutions to the problem of upgrading the original 10 Mbps Ethernet standard to operate at 10 times that speed. They both use exactly the same packet format as 10-BaseT Ethernet, so they can be transparently bridged to 10-BaseT Ethernet.

100-BaseT Standards

All three variants of 100-BaseT are defined by IEEE 802.3 in the IEEE P802.3u/D2 Supplement.

Design Principles

100-BaseT is the closest to being conventional Ethernet running 10 times faster. It uses exactly the same CSMA/CD (*Carrier Sense Media Access with Collision Detection*) method as 10-BaseT for sharing the physical network medium between devices.

Media Access

When a CSMA/CD device has a data packet to transmit, it first waits for a period when the network medium (the wire) is quiet, then starts to transmit its data while continuing to monitor the wire. If it detects that one or more other devices are simultaneously transmitting data, it stops transmitting and registers a *collision.* All the other devices that were attempting to transmit data do the same. All the devices with data still to transmit wait for a short random time before attempting to retransmit. The waiting period for each device that has registered a collision must be random, otherwise the devices would continue to retransmit at exactly the same time as each other.

The CSMA/CD protocol imposes a strict limit on the maximum size of a network segment (usually referred to as a *collision domain*). Every device must be able to detect the transmissions from every other device within a certain maximum time so that collisions can be reliably recorded. A normal 10 Mbps Ethernet network can have a maximum distance of about 2000 metres between the furthest apart UTP devices. Since 100-BaseT runs at 10 times the speed, the maximum distance is reduced to about 200 metres. Repeaters introduce further delays to the propagation of signals over the network, and 100-BaseT is allowed a maximum of only two. Large 100-BaseT networks need to be constructed from multiple small collision domains by connecting them with bridges or switches.

Physical implementation

As with 10 Mbps Ethernet, 100-BaseT has been designed to use several different physical media types:

- *100-BaseTX:* 2 pairs from a category 5 UTP cable
- *100-BaseT4:* 4 pairs from a category 3 or better UTP cable
- *100-BaseFX:* Optical fiber.

Nearly all the 100-BaseT equipment that is currently available uses the 100-BaseTX technology. 100-BaseTX requires a cabling installation in which all cable runs and connectors are properly installed to EIA/TIA category 5 standards and in which end-to-end cable runs are less than 100 metres.

100-BaseTX physical details

100-BaseTX uses the data encoding and transmission parts of TP-PMD (formerly known as CDDI), an ANSI standard originally developed for running FDDI over UTP. The data to be transmitted is chopped into 4-bit sequences which are each encoded into 5-bit quintets (4B5B encoding). The 5-bit quintets are then scrambled in order to produce a more random bit order.

continued on next page

100VG and 100-BaseT *(continued)*

Finally, the 5-bit quintets are transmitted sequentially on a single pair of wires using the MLT-3 encoding scheme. MLT-3 converts the binary digits (0s and 1s) physically into one of three voltage levels applied to the wires; –1, 0, and +1 volts. For every binary 1 transmitted, the voltage changes to the next level in the sequence –1, 0, +1, 0, –1..., while for every 0 transmitted, the voltage stays at its existing level. MLT-3 reduces the average frequency of the signal on the wires and thus reduces its susceptibility to interference from external electromagnetic radiation.

100 VG Standards

100VG is defined by the IEEE 802.12 standard.

Design principles

In contrast to 100-BaseT, 100VG was designed as a direct replacement for 10-BaseT Ethernet that could run on the same category 3, 4, or 5 *Unshielded Twisted Pair* (UTP) cabling infrastructure in the same configuration—VG stands for "voice grade" and refers to the category 3 cable that can be used. It uses a completely different method from Ethernet—*Demand Priority*—to share the network medium between devices. Up to 3 levels of 100VG hub can be cascaded to construct single large local area network segments covering the same physical area as existing 10-BaseT networks.

100VG can also operate with Token Ring frames, in which case it becomes a direct replacement for existing Token Ring networks. The Token Ring and Ethernet frame types cannot be mixed on the same network segment; two such networks must be joined by a translating bridge or router in the same way as conventional Token Ring and Ethernet.

Media Access

Network access using the Demand Priority protocol is controlled by the hubs. A device that needs to transmit data issues a request to the hub. The hub continuously scans only those devices with data to transmit and lets each of them transmit one packet in turn. The packet is routed by the hub to its destination port by looking up its destination address in an internal table listing the addresses of all connected devices.

Where several hubs are connected together, one of them becomes the level 1 or root hub and each hub separately continues to scan its connected devices. A hub that needs to send data to a device connected to a different hub raises a request with the next hub in the chain that includes the number of devices it has detected with pending requests to send. It is allowed to send one packet for each such device in a single transmission.

Two priority levels—*normal* and *high*—are defined for 100VG packets. The high-priority status is intended for time-critical applications, such as some multimedia implementations. High priority packets are always serviced first, with the proviso that if excessive delays develop in servicing normal priority packets, they have their status raised to high so that service for normal packets never completely ceases. The priority system can only be used by applications that explicitly take account of its availability and have been written accordingly.

Physical implementation

100VG as currently implemented operates using all four pairs from a standard category 3, 4, or 5 UTP cable, as defined in the EIA/TIA 568 wiring standard. Versions using fiber optic cable and 2 pairs of UTP or STP are also planned.

For transmission over the 100VG physical medium, data, originally in 8-bit bytes (called *octets* in a communications context), is treated as a stream of bits. The bit stream is divided into 5-bit quintets which are sequentially distributed between four communications channels, representing the four pairs of wires. Before transmission, each quintet is scrambled according to a different set of rules for each of the four channels with the intention of randomizing the bit pattern on each channel— random bit patterns are less susceptible to radio frequency interference and cause less transfer of signals (*crosstalk*) between adjacent pairs of wire. The scrambled quintets are now mapped into predetermined 6-bit sequences (symbols) that contain equal numbers of 0s and 1s; this is called 5B6B encoding. There are 32 (2 raised to the power 5) possible 5-bit sequences and only 20 available balanced 6-bit symbols, so 12 of the 5-bit sequences are encoded by two unbalanced but complementary 6-bit symbols each, used alternately. The use of a small number of balanced symbols ensures that there will always be equal numbers of 0s and 1s transmitted on the wires—important for the correct operation of the receiving equipment—and aids error checking; a symbol with more than 3 consecutive 0s or 1s is always invalid.

Finally, the 6-bit symbols encoding the original data are transmitted simultaneously over the 4 pairs of wires using a 30MHz clock on each pair and *non-return-to-zero* (NRZ) encoding. NRZ encoding allows one bit to be transmitted per clock cycle; a high voltage represents a 1, a low voltage a 0. Under NRZ, if two consecutive bits are both 1s or 0s, the voltage does not change as they are received. However, the 5B6B encoding ensures that there can never be more than 3 identical bits in a row, and therefore that the voltage will change regularly, which is important for keeping the receiving circuits in synchronisation with the sending ones.

Key differences

We now compare some key aspects of 100VG and 100-BaseT.

Cable

100-BaseTX requires 2 wire pairs of a fully category 5 compliant cabling, with cable runs of less than 100 metres. Some installations that use Category 5 compliant cable will not meet this requirement because of poor installation or the use of non-Category 5 connectors.

100VG will operate over 100 metres of Category 3 or 4, or 150 metres of Category 5 cable. However, it requires all 4 wire pairs in the cable and these will not always be available, either because only 2 pairs were installed or because some of the pairs are already in use for other services.

Hubs

The 100-BaseT standard describes two classes of hub. Type I hubs can interconvert signals between different physical media types (e.g., 100-BaseTX and 100-BaseT4); Type II hubs just repeat the signal. There can only be one Type I hub in a 100-BaseT domain, whereas there can be one or two Type II hubs. The lead connecting the two hubs must be 5 metres or less in length.

Up to three levels of 100VG hub can be interconnected to form networks containing very many hubs—in general, 10-BaseT repeaters can be directly replaced by 100VG hubs.

ROGER COHEN holds a science degree from Cambridge University and spent 4 years doing postgraduate research in medical biochemistry. After careers in lexicography and text processing, he became network manager for various major UK companies. Since 1989 he has been continuously developing a test suite for assessing the performance of network components, technologies and operating systems, as well as designing and troubleshooting large LANs and WANs. He is now technical director of Netmarq, a specialist network test house, and also of Homepage, a Web design and service provision company. E-mail: `netmarq@cix.compulink.co.uk`

Call for Papers

The ACM *SIGCOMM '97* conference, will be held September 14–18 1997 at the Palais des Festivals, Cannes, France. SIGCOMM '97 seeks papers about significant contributions to the broad field of computer and data communication networks.

Topics

Authors are invited to submit full papers concerned with both theory and practice. The areas of interest include, but are not limited to:

- Analysis/design of computer network architectures/algorithms
- Scalable architectures and algorithms for large networks
- Resource sharing and quality of service in networks
- Network support for multimedia
- Experimental results from operational networks
- High-speed networks
- Routing and addressing
- Wireless networking, support for mobile hosts
- Distributed application infrastructure paradigms
- Distributed common application services, middleware protocols
- Network management
- Protocol specification, verification, and analysis

Format

SIGCOMM '97 is a single-track, highly selective conference. Successful submissions typically report results firmly substantiated by experiment, implementation, simulation, or mathematical analysis; however, more qualitative explorations of important architectural issues are also encouraged.

In addition to the presentation of papers and results, SIGCOMM '97 will offer tutorials by noted instructors on the two days preceding the actual conference.

Submissions

Papers must be less than 20 double-spaced pages long (or 12 pages in camera-ready format), have an abstract of 100–150 words, and be original material presenting ideas and results that have not been previously published nor are currently under review by another conference or journal. Any previous or simultaneous publication of related material should be explicitly noted in the submission.

Important dates

Paper submissions:	31 January 1997
Tutorial proposals:	31 January 1997
Notification of acceptance:	28 April 1997
Camera ready papers due:	30 May 1997

Due to the high number of anticipated submissions, authors are encouraged to strictly adhere to the submission date. Papers will not be accepted after the paper submission deadline unless an extension has been granted by the Program Co-Chairs.

All submitted papers will be judged based on their quality and relevance through double-blind reviewing where the identities of the authors are withheld from the reviewers. Authors' names should not appear on the paper or in the *PostScript* file for electronic submissions. Authors of accepted papers will need to sign an ACM copyright release form and present their paper at the conference. The Proceedings of the conference will be published as a special issue of ACM SIGCOMM *Computer Communication Review*. The program committee may also select a few papers for possible publication in the IEEE/ACM *Transactions on Networking*.

Submission Guidelines

Papers must be submitted electronically. Instructions for submission will be described on our web page:

`http://www.inria.fr/rodeo/sigcomm97/submit.html`

Tutorials

SIGCOMM '97 will begin with two days of tutorials, each of which is intended to cover a single topic in detail. Proposals are solicited from individuals willing to give tutorials, which should be a full day in length and cover topics at an introductory or advanced level. Tutorial submissions should be made to the Tutorial Chair noted below and include an extended abstract and outline (2–4 pages), and an indication of length, objectives, and intended audience.

Student Paper Award

Papers submitted by students will enter a student-paper award contest. Among the accepted papers, a maximum of four outstanding papers will be awarded full conference registration and a travel grant of $800 US dollars. To be eligible the student must be the sole author of the paper, or the first author and primary contributor. A cover letter must identify the paper as a candidate for this competition.

SIGCOMM Award

The keynote speaker at SIGCOMM '97 will be the 1997 winner of the ACM SIGCOMM Award for lifetime contributions to the field of computer communication. Procedures for nominating candidates for the SIGCOMM Award can be obtained from David C. Wood (wood@mitre.org).

Conference committee

Program Co-Chairs:

Christian Huitema
Bellcore, MCC 1J236B
445 South Street
Morristown, NJ 07960-6438
U.S.A.
Ph: +1 (201) 829 4266
Fax: +1 (201) 829 2504
huitema@bellcore.com

Scott Shenker
Xerox PARC
3333 Coyote Hill Road
Palo Alto, CA 94304–1314
U.S.A.
Ph: +1 (415) 812 4840
Fax: +1 (415) 812 4471
shenker@parc.xerox.com

General Chair:

Christophe Diot
INRIA
2004 Route des Lucioles
BP 93
06902 Sophia Antipolis
FRANCE
Ph: +33 4 93 65 78 25
Fax: +33 4 93 65 77 65
christophe.diot@sophia.inria.fr

Tutorial Chair:

Walid Dabbous
INRIA
2004 Route des Lucioles
BP 93
06902 Sophia Antipolis
FRANCE
Ph: +33 4 93 65 78 25
Fax: +33 4 93 65 77 65
walid.dabbous@sophia.inria.fr

Program Committee

M. Baker, Stanford Univ., USA; J. Crowcroft, UCL, UK; A. Danthine, Univ. de Liege, Belgium; P. Danzig, USC, USA; C. Diot, INRIA, France; P. Druschel, Rice Univ., USA; D. Estrin, USC, USA; S. Floyd, LBL, USA; P. Francis, NTT, Japan; J. J. Garcia-Luna-Aceves, UCSC, USA; R. Guerin, IBM, USA; M. Handley, ISI, USA; C. Huitema, Bellcore, USA; P. Humblet, Eurecom, France; S. Keshav, Cornell Univ., USA; J. Kurose, Univ. of Massachusetts, USA; S. Lam, Univ. of Texas, USA; W. Leland, Bellcore, USA; S. McCanne, U. C. Berkeley, USA; G. Neufeld, UBC, Canada; D. Oran, Cisco, USA; C. Partridge, BBN, USA; V. Paxson, LBL, USA; S. Pink, SICS, Sweden; D. Schmidt, Washington Univ., USA; S. Shenker, Xerox PARC, USA; K. Sollins, MIT, USA; J. Turner, Washington Univ., USA; M. Steenstrup, BBN, USA; G. Varghese, Washington Univ., USA; H. Zhang, CMU, USA; L. Zhang, UCLA, USA.

Call for Papers

Topics

The third *IEEE ATM Workshop* will be held May, 26–28, 1997 in Lisboa, Portugal. The workshop is sponsored by the IEEE Communications Society and has the objective of fostering the exchange of information among those working in the diverse areas of ATM. The Workshop Committee invites submissions and participation from researchers and developers in academia, industry and government. The Workshop is intended for those who are actively involved in or following ATM research and development. Original contributions are invited on topics such as:

- ATM switch architectures, implementation and performance
- Admission control and routing
- ATM internetworking
- Scheduling for service integration (e.g., Weighted Fair Queueing)
- Wireless ATM (terrestrial and satellite based)
- Experience with operational ATM networks
- ATM LAN switches
- ATM field trials
- ATM traffic modelling
- IP/ATM interworking (TCP over ATM)
- Video/image coding and transmission over ATM
- Telephony over ATM
- ATM network design (access network, core network)
- ATM standards
- Broadband signalling
- Network management and operation
- Traffic management functions and procedures
- Multicasting over ATM
- Mapping services to transfer capabilities (ABR, VBR/SBR...)

Important dates

Submission of extended abstract (5 pages):	January 31, 1997
Authors notified:	March 15, 1997
Full paper due (10 pages max):	April 15, 1997.

Submissions

Submit by e-mail (*PostScript*), fax or regular mail, your extended abstract to:

Jim Roberts
CNET/PAA/ATR
38, rue du General-Leclerc
92131 Issy les Moulineuax
FRANCE
Tel: +33 1 4529 5701
Fax: +33 1 4529 6069
E-mail: `james.roberts@issy.cnet.fr`

Please include complete information on author(s), affiliation(s), address(es), telephone, fax and e-mail.

More information

For registration and general information, contact:

Augusta Casaca
INESC
Rua Alves Redol, 9
1000 Lisboa
PORTUGAL
Tel: +351 1 3100233
Fax: +351 1 3145843
E-mail: `augusto.casaca@inesc.pt`

Future NetWorld+Interop Dates and Locations

NetWorld+Interop 96	Sydney, Australia	November 25–29, 1996
NetWorld+Interop 97	Singapore	April 7–11, 1997
NetWorld+Interop 97	Las Vegas, NV	May 5–9, 1997
NetWorld+Interop 97	Frankfurt, Germany	May 12–15, 1997
NetWorld+Interop 97	Tokyo, Japan	June 2–6, 1997
NetWorld+Interop 97	Atlanta, GA	October 6–10, 1997
NetWorld+Interop 97	Paris, France	October 20–23, 1997
NetWorld+Interop 97	London, England	October 27–30, 1997
NetWorld+Interop 97	Sydney, Australia	November 25–28, 1997

All dates are subject to change.

More information

Call 1-800-INTEROP or +1-415-578-6900 for more information. Or send e-mail to `info@interop.com` or fax to +1-415-525-0194. For the latest information about Interop DotCom and NetWorld+Interop as well as other SOFTBANK produced events, check our *Interop Online* home page at `http://www.interop.com`

NetWorld+Interop is produced by SOFTBANK Exposition and Conference Company, 303 Vintage Park Drive, Foster City, California 94404–1138, USA.

Write to *ConneXions*!

We'd love to hear your comments, suggestions and questions about anything you read in *ConneXions*. Our editorial address is given below. Use it for letters to the Editor, requests for the index of back issues, questions about particular articles etc.:

ConneXions—The Interoperability Report
303 Vintage Park Drive
Suite 201
Foster City
California 94404–1138
USA
Phone: +1 415-578-6900 or 1-800-INTEROP (Toll-free in the USA)
Fax: +1 415-525-0194
E-mail: `connexions@interop.com`
URL: `http://www.interop.com`

Subscription information

For questions about your subscription please call our customer service hotline: 1-800-575-5717 or +1 610-892-1959 outside the USA. This is the number for our subscription agency, Seybold Publications. Their fax number is +1 610-565-1858. The mailing address for subscription payments is: P.O. Box 976, Media, PA 19063–0976.

This publication is distributed on an "as is" basis, without warranty. Neither the publisher nor any contributor shall have any liability to any person or entity with respect to any liability, loss, or damage caused or alleged to be caused, directly or indirectly, by the information contained in *ConneXions—The Interoperability Report®*

10.11

CONNEXIONS
303 Vintage Park Drive
Suite 201
Foster City, CA 94404-1138
Phone: 415-578-6900
FAX: 415-525-0194

ADDRESS CORRECTION
REQUESTED

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

CONNEXIONS

EDITOR and PUBLISHER Ole J. Jacobsen

EDITORIAL ADVISORY BOARD

Dr. Vinton G. Cerf
Senior Vice President, MCI Telecommunications
President, The Internet Society (1992 – 1995)

A. Lyman Chapin, Chief Network Architect,
BBN Communications

Dr. David D. Clark, Senior Research Scientist,
Massachusetts Institute of Technology

Dr. David L. Mills, Professor,
University of Delaware

Dr. Jonathan B. Postel, Communications Division Director,
University of Southern California, Information Sciences Institute

Printed on recycled paper

Subscribe to CONNEXIONS

U.S./Canada ❑ $195. for 12 issues/year   All other countries ❑ $245. for 12 issues/year

Name ______________________ Title ______________________
Company ______________________ E-mail ______________________
Address ____
City ______________________ State ____________ Zip ____________
Country ______________________ Telephone (   ) ______________
Fax (   ) ______________

❑ Check enclosed (in U.S. dollars made payable to CONNEXIONS).
❑ Visa ❑ MasterCard ❑ American Express ❑ Diners Club Card# ______________ Exp.Date ________
Signature ____

Please return this application with payment to: CONNEXIONS
303 Vintage Park Drive, Suite 201
Foster City, CA 94404-1138
415-578-6900 FAX: 415-525-0194
connexions@interop.com

Back issues available upon request $15./each
Volume discounts available upon request